

Imputation Strategies Under Clinical Presence: Impact on Algorithmic Fairness

Vincent Jeanselme

*MRC Biostatistics Unit,
University of Cambridge, Cambridge, UK*

VINCENT.JEANSELME@MRC-BSU.CAM.AC.UK

Maria De-Arteaga*

*McCombs School of Business,
University of Texas at Austin, Austin, USA*

DEARTEAGA@MCCOMBS.UTEXAS.EDU

Zhe Zhang*

*Rady School of Management,
University of California, San Diego, USA*

ZHE@RADY.UCSD.EDU

Jessica Barrett

*MRC Biostatistics Unit,
University of Cambridge, Cambridge, UK*

JESSICA.BARRETT@MRC-BSU.CAM.AC.UK

Brian Tom

*MRC Biostatistics Unit,
University of Cambridge, Cambridge, UK*

BRIAN.TOM@MRC-BSU.CAM.AC.UK

Abstract

Biases have marked medical history, leading to unequal care affecting marginalised groups. The patterns of missingness in observational data often reflect these group discrepancies, but the algorithmic fairness implications of group-specific missingness are not well understood. Despite its potential impact, imputation is too often an overlooked preprocessing step. When explicitly considered, attention is placed on overall performance, ignoring how this preprocessing can reinforce group-specific inequities. Our work questions this choice by studying how imputation affects downstream algorithmic fairness. First, we provide a structured view of the relationship between clinical presence mechanisms and group-specific missingness patterns. Then, through simulations and real-world experiments, we demonstrate that the im-

putation choice influences marginalised group performance and that no imputation strategy consistently reduces disparities. Importantly, our results show that current practices may endanger health equity as similarly performing imputation strategies at the population level can affect marginalised groups differently. Finally, we propose recommendations for mitigating inequities that may stem from a neglected step of the machine learning pipeline.

Keywords: Clinical Presence, Fairness, Imputation

1. Introduction

Machine learning models for healthcare often rely on observational data. At the core of observational data generation is a complex interaction between patients and the healthcare system, which we refer to as *clinical presence* (Jeanselme et al., 2022). Each observation, from orders of laboratory tests

* Equal contribution.

to treatment decisions, reflects access to medical care, patients’ medical states, and also practitioners’ expertise and potential biases. Historically, healthcare access, treatment and outcomes have been marked by inequalities (Chen et al., 2021; Freeman and Payne, 2000; Jeanselme et al., 2021; Kim et al., 2016; Norris and Nissenon, 2008). For instance, Price-Haywood et al. (2020) hypothesised that the disproportionate mortality rate from Covid-19 among Black patients can, in part, be explained by longer waiting times before accessing care.

Clinical presence patterns can, therefore, reflect disparities. Specifically, observation and missingness can vary across groups. Developing machine learning models on these data raises ethical concerns about automating and reinforcing injustices.

Current practices for handling missing data often rely on imputing data with overall performance in mind (Emmanuel et al., 2021), without consideration of the algorithmic fairness consequences associated with this choice. Despite the risk of aggravating inequities reflected in group-specific missingness patterns, the effect of this imputation step remains understudied. In this work, we explore the impact of imputation on data imprinted by group-specific missingness patterns emerging from medical practice and historical biases. First, we identify scenarios of clinical presence that could result in group-specific missingness patterns, grounded on historical evidence of these phenomena in medicine. Then, we explore the downstream impact on group performance of standard imputation strategies on simulated data affected by this *clinical missingness*. Finally, we study group performances of different imputation strategies in real-world data.

This work provides empirical evidence that machine learning pipelines differing solely in their handling of missingness may result in distinct performance gaps between groups,

even when population performances present no difference. The choice of imputation strategy may therefore impact performance in a way that reinforces inequities against historically marginalised groups. Moreover, our experiments show that no imputation strategy consistently outperforms the others and current recommendations may harm marginalised groups. Finally, we emphasise the relevance of this analysis by providing real-world evidence of clinical missingness patterns and echo the previous results in the MIMIC III dataset.

2. Related work

This work explores the link between missingness and algorithmic fairness in machine learning for healthcare. In this section, we review related literature across domains.

2.1. Clinical missingness

Clinical missingness is a medical expression of the well-studied missingness patterns (Little and Rubin, 2019): *Missing Completely At Random* (MCAR) — random subsets of patients and/or covariates are missing, *Missing At Random* (MAR) — missing data patterns are a function of observed variables, and *Missing Not At Random* (MNAR) — missing patterns depend on unobserved variables or the missing values themselves.

Traditional statistical models are not adapted to handle missing covariates. Consequently, practitioners may rely on single imputation strategies such as mean, median, nearest neighbours (Batista et al., 2002; Bertsimas et al., 2021) or the preferred multiple imputation methods (Newgard and Lewis, 2015; Rubin, 2004; White et al., 2011). Typically, these imputation approaches assume MCAR and/or MAR patterns. They may be ill-adapted to handle informative missingness, particularly as MNAR and MAR are non-identifiable from

observational data alone and require domain expertise for adequate modelling. The recommended strategy to tackle this non-identifiability issue is to control the imputation model on additional covariates to render the MAR assumption more plausible (Haukoos and Newgard, 2007). Our work shows the potential shortcomings of this covariate-adjusted imputation strategy under group-specific missingness patterns.

2.2. Algorithmic fairness in medicine

The risk of reinforcing historical biases is of critical concern in medicine, where inequalities can have life-threatening implications. Measuring and mitigating this risk is the aim of algorithmic fairness (Chouldechova and Roth, 2020). In this paper, we follow the ‘equal performance’ group definition of algorithmic fairness (Rajkomar et al., 2018), which evaluates if the model performs comparably across groups (Chouldechova et al., 2018; Flores et al., 2016; Noriega-Campero et al., 2019).

Definition 1 (Equal Performance) *A pipeline p is fairer than another q with regard to group g if its performance gap is the smallest, i.e. $|\Delta_g(p)| < |\Delta_g(q)|$ with $\Delta_g(p) := d(p(\{X_i\}_{G_i=g})) - d(p(\{X_i\}_{G_i \neq g}))$ for some performance metric d , a pipeline p and (X_i, G_i) , the covariates and associated group for patient i .*

This metric has been leveraged to quantify models’ impact on algorithmic fairness in medicine (Chen et al., 2018, 2019; Pfohl et al., 2019; Seyyed-Kalantari et al., 2020; Zhang et al., 2020). For instance, Seyyed-Kalantari et al. (2020) demonstrates X-ray classifiers’ performance gap between marginalised groups. However, the link between imputation and algorithmic fairness has received limited attention despite the risk of clinical missingness disparities. Our work aims to fill this gap.

2.3. Algorithmic fairness and missingness

As a community, we need to understand how to best handle clinical missingness when imprinted by biases. Martínez-Plumed et al. (2019); Fricke et al. (2020) show that mean imputation presents better fairness properties compared to complete case analysis. These works focus on one imputation strategy and ignore the potential variability of the impact of different strategies. Closer to our work, Zhang and Long (2021) show that the choice of imputation may lead to different fairness gaps when enforcing synthetic missingness patterns. However, these works do not discuss how the different missingness patterns may arise in medicine, and how a specific group may be impacted differently by different imputation strategies. In our work, we study different missingness patterns that may arise as a result of the data-generating process in healthcare. Finally, Ahmad et al. (2019); Ghassemi et al. (2020); Rajkomar et al. (2018) describe multiple challenges linked to medical data, among which they state that historical biases may lead to missingness patterns that could impact fairness, but they do not empirically study this. While informative missingness has recently received revived attention (Jeanselme et al., 2022; Getzen et al., 2022), no work has studied its potential association with fairness. Our work aims to address these gaps in the literature by demonstrating the existence of this problem, characterising different types of group-specific missingness patterns in medicine, and exploring the impact of different imputation strategies under different clinical presence scenarios. In addition to showing the impact of imputation choice on fairness gaps, we highlight that the *same imputation strategy* may benefit a group under one missingness pattern but hurt this same group in

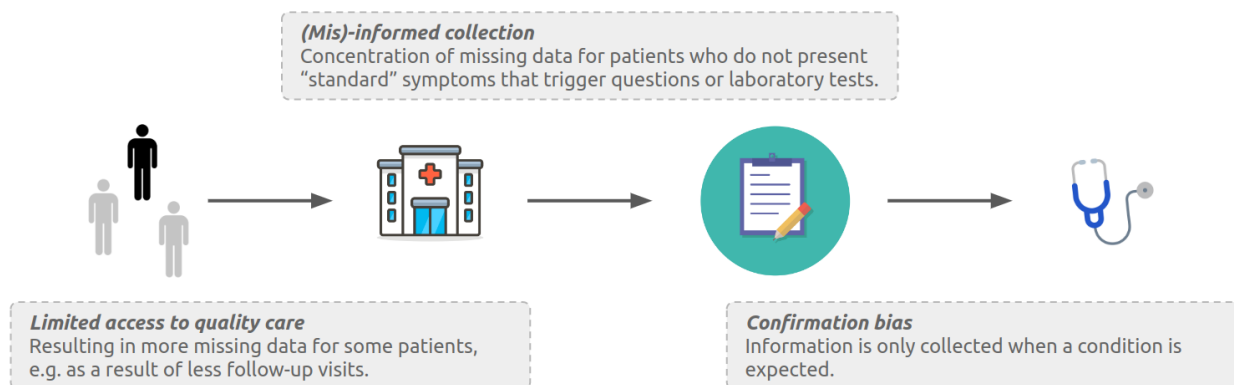


Figure 1: Examples of group-specific clinical presence mechanisms.

another. Importantly, we also show that a given group may benefit under one imputation and suffer under another imputation in the same setting, even if the two strategies perform identically at the population level. These are novel findings that invite practitioners to perform careful sensitivity analysis of imputation choice on fairness gaps.

3. Clinical missingness scenarios

This section shows how group-specific missingness can result from clinical presence. Figure 1 introduces the following scenarios:

Limited access to quality care (S1). When certain groups do not have access to the same health services, this results in more missing covariates for these groups.

Socioeconomic factors resulting from structural injustices (Barik and Thorat, 2015; Nelson, 2002; Szczepura, 2005; Yearby, 2018) such as insurance, work schedule flexibility, distance to hospitals (Barik and Thorat, 2015) or mobility, result in inconsistent medical history (Gianfrancesco et al., 2018), additional waiting time before looking for care (Weissman et al., 1991), avoidance of preventing care (Smith et al., 2018), and limited access to advanced diagnostic tools (Lin et al., 2019). This

diminished access to care is potentially reflected as missing data. For instance, patients may have no annual checkup data if their insurance does not cover or encourage this service.

(Mis)-informed collection (S2). Often, medical research has focused on a subset of the population. The resulting guidelines may be ill-adapted to other groups and relevant covariates may be missing due to standard recommendations.

Historically researchers focused on (perceived) highest-risk groups: breast cancer predominantly studied in women (Arnould et al., 2006; Giordano, 2018), cardiovascular disease in men (Vogel et al., 2021), skin cancers in whiter skins (Gloster Jr and Neal, 2006), and autism in men (Gould and Ashton-Smith, 2011). Resultant medical practices and guidelines target these groups. However, substantial evidence shows the prevalence of these diseases among other groups. Stemming from biological differences, different groups may present different symptoms and expressions for the same condition. The difference in disease expression and the absence of adapted tests result in missing covariates necessary to identify the disease. For instance, screening recommendations may only be prescribed conditioned

on observation of “standard” symptoms. If the symptoms considered are not the expected disease expression for a marginalised subgroup, this will result in more missing screening procedures for this group.

Confirmation bias (S3). Practitioners collect data based on expertise and informative proxies that are not recorded, e.g. patient feeling unwell.

For instance, practitioners may record the value of a test only if they suspect it will be abnormal. The literature presents evidence of this phenomenon where the presence of a specific medical test is more informative of the outcome than the test result itself (Agniel et al., 2018; Sisk et al., 2020). Wells et al. (2013) also suggest that missing laboratory tests correspond to healthy results, e.g. doctors do not collect or record data if they are irrelevant. Similarly, sicker patients present more complete data (Rusanov et al., 2014; Sharafoddini et al., 2019; Weiskopf et al., 2013).

Formalisation. Consider two covariates (X_1, X_2) influenced by the underlying condition Y and the group membership G . Note that the disease prevalence may also depend on G . One covariate X_1 is observed for all patients, while X_2 is potentially missing. Following the notations from Mohan and Pearl (2021), let O_2 be the indicator of observation of X_2 such that the observed value is defined as:

$$X_2^* = \begin{cases} \emptyset & \text{if } O_2 = 0 \\ X_2 & \text{otherwise} \end{cases}$$

In (S1), G informs O_2 because of group socioeconomic differences. In (S2) and (S3), G impacts the observation process through group-specific disease expression. While the influence of medical covariates on the missingness patterns characterises both (S2) and (S3), (S2) describes how guidelines may depend on observed covariates, whereas (S3)

reflects how the observation process may depend on X_2 itself or unobservable covariates correlated with X_2 . For instance, (S2) may consist of a guideline recommending to measure X_2 if X_1 is within a given range. However, if a patient is a member of a group for which X_1 is not informative—or for which the informative range is different— X_2 might not be observed as X_1 is not in the guideline test-triggering range. This may lead to more missing data for X_2 in the group with different characteristics for X_1 . (S3) differs as practitioners would *record* the value of X_2 only if this one is abnormal.

These dependencies result in three distinct patterns between missingness, group and covariates, summarised with directed acyclic graphs (DAGs) in Figure 2.

4. Experiments

In this section, we explore how the choice of imputation affects group-specific performance, and potentially reinforces disparities in data marked by clinical missingness. We first present simulation studies in which we enforce specific missingness patterns. This analysis allows us to control clinical missingness patterns and measure the potential impact of imputation on algorithmic fairness. We accompany these results with real-world evidence of group-specific missingness patterns and show the impact of different imputation strategies on marginalised group performance. For reproducibility, all experiments’ code is available on Github¹.

4.1. Datasets

Assume a population of N patients with associated covariates X , marginalised group membership G , and outcome of interest Y .

1. <https://github.com/Jeanselme/ClinicalPresenceFairness>




Figure 2: Directed Acyclic Graphs (DAGs) associated with the identified clinical missingness scenarios. Full circled covariates are observed, dotted ones unobserved. Y is the condition, G , the group membership, X_1 and X_2 the two covariates. O_2 is the observation process associated to X_2 . Red dependencies underline the differences between scenarios.

Simulation. We introduce a bidimensional ($X \in \mathbb{R}^2$) synthetic population ($N = 10, 100$) divided into two groups ($G \in \{0, 1\}$), and assume the marginalised group is a minority in the population with ratio 1:100. These groups differ in disease expression, i.e. positive cases across groups differ in how they express the disease. Then clinical missingness patterns are enforced on the second dimension X_2 following the scenarios introduced in Section 3. Figure 3 provides a graphical summary of how clinical missingness is enforced on the synthetic data. The associated predictive task is to classify between positives and negatives. (See Appendix A.1 for full data generation protocol reflecting the enforcement of the previously-introduced scenarios).

MIMIC III. The real-world analysis relies on the laboratory tests from Medical In-

formation Mart for Intensive Care (MIMIC III) dataset (Johnson et al., 2016). Following data harmonisation (Wang et al., 2020), we select adults who survived 24 hours or more after admission to the intensive care unit, resulting in a set of 36,296 patients sharing 67 laboratory tests. The goal is to predict short-term survival (7 days after the observation period — Y) using the most recent value of each laboratory test observed in the first 24 hours of observation (X). We select short-term survival as it is a standard task in the machine learning literature (Jeanselme et al., 2022; Nagpal et al., 2021; Tsiklidis et al., 2022; Xu et al., 2019) and the associated labels are less likely to suffer from group-specific misdiagnosis, and, therefore, disentangles our analysis from potential biases in labelling (Chen et al., 2020). In practice, deploying this model could be used for




Figure 3: Graphical summary of clinical missingness enforcement in the simulation experiments. Note that our simulations’ choices result in missingness in the marginalised group only in (S1) and (S2), but in the majority only in (S3).

care prioritisation of patients with predicted elevated risk.

4.2. Handling missing data

The simulation and MIMIC III datasets present missing data that are traditionally imputed for analysis. We consider the following common imputation strategies:

Single median imputation (Median).

Missing data are replaced by the population median of each covariate. Due to its straightforward implementation, this methodology remains predominant in the literature despite known shortcomings (Rubin, 1976; Sinharay et al., 2001; Crawford et al., 1995).

Multiple Imputation using Chained Equation (MICE).

Missing data are iteratively drawn from a regression model built over all other available covariates after median initialisation. This approach is repeated I times with an associated predictive model for each imputed draw. At test time, the same imputation models generate I imputed points for which models’ predictions are averaged. MICE is recommended in the literature (Janssen et al., 2010; Newgard and Haukoos, 2007; Wood et al., 2004; Zhou

et al., 2001; White et al., 2011) as it quantifies the uncertainty associated with missingness. In the experiments, we used 10 iterations repeated 10 times resulting in $I = 10$ datasets with associated predictive models.

Group MICE. The previous MICE methodology assumes a MAR mechanism. To make this assumption more plausible, Haukoos and Newgard (2007) recommend the addition of potentially informative covariates. In our experiment, we, therefore, rely on both group membership and covariates for imputing the missing data ($\tilde{X} \sim X, G$ with \tilde{X} representing the imputed covariates).

Group MICE Missing. Encoding missingness has been shown to improve performance when the patterns of missingness are informative (Groenwold, 2020; Lipton et al., 2016; Saar-Tsechansky and Provost, 2007; Sperrin et al., 2020). As clinical missingness can contain informative patterns (Jeanselme et al., 2022; Lipton et al., 2016), we concatenate missingness indicators to the imputed data from Group MICE (Appendix A explores the concatenation of missing indicators with the other strategies).




Figure 4: AUC performance gaps Δ and group-specific AUCs across scenarios on 100 synthetic experiments. If $\Delta < 0$, the marginalised group has worse AUC than the majority.

4.3. Experimental setting

After imputation, each pipeline relies on a logistic regression model — a pillar in medicine (Nick and Campbell, 2007; Goldstein et al., 2017) — to discriminate between positive and negative cases ($Y \sim \tilde{X}$).

Adopting the *equal performance across groups* definition (Rajkomar et al., 2018) of algorithmic fairness, we measure each pipeline’s discriminative performances for the different groups. We use the Area Under the Curve for the Receiver Operating Characteristic curve (AUC - ROC, i.e. d in Section 2.2) as proposed in Rööslı et al. (2022); Larrazabal et al. (2020); Zhang et al. (2022).

This metric quantifies algorithmic fairness but does not quantify how deployment can hurt subgroups at a fixed threshold on the predicted risk. In the MIMIC III study, we measure the False Negative Rate (FNR) assuming the availability of priority care for 30% of the population (sensitivity to this threshold is presented in Appendix A.2). In the 30% highest-risk population, we measure the prioritisation — the group-specific proportion of patients who would receive care under this policy — and misclassification rates in the groups of interest. In this setting, FNR corresponds to the non-prioritisation of high-risk patients. The gap in FNR be-

tween groups answers the question: how marginalised groups would be incorrectly deprioritized? Additional experimental design descriptions and results are provided in Appendix A.

5. Results

This section presents the insights obtained through both simulations and real-world experiments.

5.1. Simulations

We conduct 100 simulations in which the three clinical presence scenarios are independently enforced. We apply the imputation strategies described in Section 4.2 and train a logistic regression with l2 penalty ($\lambda = 1$). Results are computed on a 20% test set and averaged over the 100 simulations. Figure 4 presents the AUC gap (Δ defined in Section 2.2) between the majority and the minority, and group-specific AUCs.

Insight 1: Equally-performing imputation strategies at the population level can result in different marginalised group performances. Consider (S1), all imputation methodologies result in similar population AUCs, as shown by the grey dots. However, note how the AUC evaluated on

the marginalised group presents a gap of 0.1 between MICE and Group MICE. This phenomenon is explained by how imputation strategies result in different imputed covariate distributions. The logistic regressions built on these imputed data would weigh covariates differently and then have different predicted values.

Insight 2: No strategy consistently outperforms the others across clinical presence scenarios. Population-level performances remain stable between Group MICE and MICE over all scenarios, but these strategies have contrasting marginalised group AUCs. Importantly, Group MICE should be preferred in (S1) as it minimises the performance gap. For the same reason, MICE should be used in (S2), whereas both methodologies present inconclusive fairness differences in (S3). While this result is specific to this simulation, this exemplifies how no methodology consistently reduces the performance gap across groups.

Insight 3: Current recommendation of leveraging additional covariates to satisfy MAR assumption, or using missingness indicators can harm marginalised group’s performance. Note how Group MICE presents *worse* performance than MICE in (S2). The recommendation of including additional covariates to make the MAR assumption more plausible is not always suitable as it may add *noise* and lead to poorer performance. In another example, see how the model considering missingness provides an edge in (S3) compared to Group MICE but hurts performance in (S1). This observation reinforces the necessity of measuring the performance sensitivity to imputation. Additionally, it underlines how understanding the missingness process is essential to control for relevant covariates.

5.2. MIMIC III

In this real-world experiment, we consider groups defined by the following attributes: ethnicity (Black vs non-Black), sex (female vs male), and insurance (publicly vs privately insured). Table 1 shows the number of orders and the number of distinct laboratory tests (out of the 67 possible tests) performed during the first-day post-admission for each subgroup. This last number reflects the missingness of the vector used for prediction.

For this experiment, patients are split into three sets: 80% for training, 10% for hyperparameter tuning and 10% for testing. We perform a 12 penalty search for the logistic regression among $\lambda \in [0.1, 1, 10, 100]$. Table 2 presents predictive performances at the population level averaged on the bootstrapped test set over 100 iterations. Assuming capacity for additional care for the 30% highest risk, we explore care prioritisation. Figure 5 displays our main results: the gaps in prioritisation and the false negative rates stratified by groups of interest under the different imputation strategies.

Insight 4: Real-world data presents group-specific clinical presence patterns. While the causes of clinical missingness *cannot* be distinguished from observational data alone, one can observe evidence of non-random missingness patterns in the MIMIC III dataset, as shown in Table 1. Specifically, note the larger number of orders for patients who die during their stay compared with the ones who survive. This pattern is consistent with a possible *confirmation bias* scenario (S3), if doctors are monitoring sicker patients more closely. Another example of non-random missingness is that there are fewer test orders for female, Black, and publicly insured patients, but little difference in the diversity of tests prescribed. While this may be explained by the underlying conditions or other medically




Figure 5: Prioritisation performance gaps Δ across marginalised groups in MIMIC III experiment. If $\Delta > 0$, the marginalised group has a larger value of the given metric than the rest of the population.

relevant factors, the combination of similar diversity of tests but less frequent observations results in a less up-to-date patient’s health status for modelling. Thus, even though the cause of testing differences is unclear, these observations show the connection between testing patterns, group membership, and outcomes. This real-world evidence of non-random missingness patterns among subgroups of patients raises concerns about increasing inequities if the fairness implications of imputation methods are not considered.

Insight 5: Marginalised groups can benefit or be harmed by equally performing imputation strategies at the population level. Note how MICE and Group MICE perform similarly at the population level in Table 2, but present different performances for marginalised groups (see Figure 5). Consider the ethnicity split: these methodologies have opposite consequences on Black patients. MICE would result in more care for Black patients and a smaller gap in FNR. By contrast, Group MICE would halve prioritisation and double the FNR gap in favour of non-Black patients. Crucially, this difference solely results from the imputation strategy adopted in these two pipelines.

Table 1: Mean (std) number of orders and observed tests performed during the first post-admission stratified by marginalised group and outcomes.

	Orders		Distinct tests	
Alive ⁺	5.68 (4.64)	*	40.80 (6.73)	*
Dead ⁺	7.57 (5.44)		37.22 (7.50)	
Black	5.24 (4.08)	*	40.94 (6.94)	*
Other	5.86 (4.77)		40.52 (6.84)	
Female	5.54 (4.45)	*	40.75 (6.89)	*
Male	6.03 (4.91)		40.41 (6.80)	
Public	5.67 (4.57)	*	40.46 (6.76)	*
Private	6.11 (5.01)		40.75 (7.01)	

⁺ By the 8th day after admission.

* Significant t-test p-value (< 0.001).

Table 2: Predictive performance under different imputation strategies. Mean (std) computed on the test set bootstrapped 100 times.

	AUC ROC
Group MICE Missing	0.786 (0.009)
Group MICE	0.738 (0.012)
MICE	0.742 (0.012)
Median	0.748 (0.011)

Insight 6: Different marginalised groups may be impacted oppositely by the same imputation strategy.

Female and publicly insured patients have higher prioritisation rates under all imputation methods. However, these groups show opposite gaps in their FNR compared to their counterparts (men and privately insured patients): women have more false negative cases missed while those publicly insured have fewer false negatives.

In another case of opposite impacts of imputation, Group MICE presents the smallest FNR performance gap for sex, but the largest gaps for both ethnicity and insurance. Group MICE also results in better FNR performance for publicly insured but worse for Black patients. This observation underlines the importance of identifying marginalised groups in development and deployment populations. The optimal trade-off between group and population performances, and between marginalised groups, needs to be considered as different pipelines could have opposite impacts.

6. Discussion

This paper is motivated by how interactions between patients and the healthcare system can result in group-specific missingness patterns. We show that resultant inequities in clinical missingness can impact downstream algorithmic fairness under different imputation strategies. This analysis demonstrates that no imputation strategy consistently provides better performances for marginalised groups. In particular, a model providing an edge in one setting can underperform in another, or even harm a different group. Moreover, the experiments conducted using the MIMIC-III dataset demonstrate the relevance of the identified problem as more than a merely theoretical concern, showing that it

is present in a widely used electronic health record dataset.

Note that our work does not claim that the specific patterns we observe will necessarily be present in other datasets. As we have emphasised, different combinations of missingness processes may lead to different fairness gaps and interactions between imputation and group performance. It may even lead to equal fairness performance of all imputation strategies, but one cannot know this a priori.

Learning from medical data without sufficient attention to the potential entanglement of clinical missingness and historical biases could reinforce and automatise inequities, and further harm historically marginalised groups. This work calls for caution in the use of imputation to reach health equity. We invite practitioners to:

- Record protected attributes and identify marginalised groups.
- Explore the practitioner-patient interaction process to identify clinical missingness disparities.
- Report the assumptions made at each stage of the pipeline.
- Perform sensitivity analysis on imputation to understand its impact on algorithmic fairness.

Future work will theoretically define in which settings the presented results stand and how model choice could mitigate discrepancies in the missingness patterns. Moreover, clinical missingness is only one dimension of how clinical presence shapes the data-generating process. The temporality and irregularity of medical time series may convey group-specific disparities that machine learning methods may amplify.

Acknowledgments

The authors would like to thank Changjian Shui (McGill University) for constructive feedback on the manuscript. This work has been partially funded by UKRI Medical Research Council (MC_UU_00002/5 and MC_UU_00002/2) and the NIH through grant R01NS124642.

References

- Denis Agniel, Isaac S Kohane, and Griffin M Weber. Biases in electronic health record data due to processes within the health-care system: retrospective observational study. *Bmj*, 361, 2018.
- Muhammad Aurangzeb Ahmad, Carly Eckert, and Ankur Teredesai. The challenge of imputation in explainable artificial intelligence models. *arXiv preprint arXiv:1907.12669*, 2019.
- N Arnould, O Pouget, M Gharbi, and JP Brettes. Breast cancer in men: are there similarities with breast cancer in women? *Gynecologie, Obstetrique & Fertilité*, 34(5):413–419, 2006.
- Debasis Barik and Amit Thorat. Issues of unequal access to public health in india. *Frontiers in public health*, 3:245, 2015.
- Gustavo EAPA Batista, Maria Carolina Monard, et al. A study of k-nearest neighbour as an imputation method. *HIS*, 87 (251-260):48, 2002.
- Dimitris Bertsimas, Agni Orfanoudaki, and Colin Pawlowski. Imputation of clinical covariates in time series. *Machine Learning*, 110(1):185–248, 2021.
- Irene Chen, Fredrik D Johansson, and David Sontag. Why is my classifier discriminatory? *Advances in Neural Information Processing Systems*, 31, 2018.
- Irene Y Chen, Peter Szolovits, and Marzyeh Ghassemi. Can ai help reduce disparities in general medical and mental health care? *AMA journal of ethics*, 21(2):167–179, 2019.
- Irene Y Chen, Emma Pierson, Sherri Rose, Shalmali Joshi, Kadija Ferryman, and Marzyeh Ghassemi. Ethical machine learning in healthcare. *Annual Review of Biomedical Data Science*, 4, 2020.
- Richard J Chen, Tiffany Y Chen, Jana Lipkova, Judy J Wang, Drew FK Williamson, Ming Y Lu, Sharifa Sahai, and Faisal Mahmood. Algorithm fairness in ai for medicine and healthcare. *arXiv preprint arXiv:2110.00603*, 2021.
- Alexandra Chouldechova and Aaron Roth. A snapshot of the frontiers of fairness in machine learning. *Communications of the ACM*, 63(5):82–89, 2020.
- Alexandra Chouldechova, Diana Benavides-Prado, Oleksandr Fialko, and Rhema Vaithianathan. A case study of algorithm-assisted decision making in child maltreatment hotline screening decisions. In *Conference on Fairness, Accountability and Transparency*, pages 134–148. PMLR, 2018.
- Sybil L Crawford, Sharon L Tennstedt, and John B McKinlay. A comparison of analytic methods for non-random missingness of outcome data. *Journal of clinical epidemiology*, 48(2):209–219, 1995.
- Tlanelo Emmanuel, Thabiso Maupong, Dimane Mpoeleng, Thabo Semong, Banyatsang Mphago, and Oteng Tabona. A survey on missing data in machine learning. *Journal of Big Data*, 8(1):1–37, 2021.
- Anthony W Flores, Kristin Bechtel, and Christopher T Lowenkamp. False positives, false negatives, and false analyses:

- A rejoinder to machine bias: There’s software used across the country to predict future criminals. and it’s biased against blacks. *Fed. Probation*, 80:38, 2016.
- Harold P Freeman and Richard Payne. Racial injustice in health care, 2000.
- Christian Fricke et al. Missing fairness: The discriminatory effect of missing values in datasets on fairness in machine learning. 2020.
- Emily Getzen, Lyle Ungar, Danielle Mowery, Xiaoqian Jiang, and Qi Long. Mining for equitable health: Assessing the impact of missing data in electronic health records. *medRxiv*, 2022.
- Marzyeh Ghassemi, Tristan Naumann, Peter Schulam, Andrew L Beam, Irene Y Chen, and Rajesh Ranganath. A review of challenges and opportunities in machine learning for health. *AMIA Summits on Translational Science Proceedings*, 2020:191, 2020.
- Milena A Gianfrancesco, Suzanne Tamang, Jinoos Yazdany, and Gabriela Schmajuk. Potential biases in machine learning algorithms using electronic health record data. *JAMA internal medicine*, 178(11):1544–1547, 2018.
- Sharon H Giordano. Breast cancer in men. *New England Journal of Medicine*, 378(24):2311–2320, 2018.
- Hugh M Gloster Jr and Kenneth Neal. Skin cancer in skin of color. *Journal of the American Academy of Dermatology*, 55(5):741–760, 2006.
- Benjamin A Goldstein, Ann Marie Navar, Michael J Pencina, and John Ioannidis. Opportunities and challenges in developing risk prediction models with electronic health records data: a systematic review. *Journal of the American Medical Informatics Association*, 24(1):198–208, 2017.
- Judith Gould and Jacqui Ashton-Smith. Missed diagnosis or misdiagnosis? girls and women on the autism spectrum. *Good Autism Practice (GAP)*, 12(1):34–41, 2011.
- Rolf HH Groenwold. Informative missingness in electronic health record systems: the curse of knowing. *Diagnostic and prognostic research*, 4(1):1–6, 2020.
- Jason S Haukoos and Craig D Newgard. Advanced statistics: missing data in clinical research—part 1: an introduction and conceptual framework. *Academic Emergency Medicine*, 14(7):662–668, 2007.
- Kristel JM Janssen, A Rogier T Donders, Frank E Harrell Jr, Yvonne Vergouwe, Qingxia Chen, Diederick E Grobbee, and Karel GM Moons. Missing covariate data in medical research: to impute is better than to ignore. *Journal of clinical epidemiology*, 63(7):721–727, 2010.
- Vincent Jeanselme, Maria De-Arteaga, Jonathan Elmer, Sarah M Perman, and Artur Dubrawski. Sex differences in post cardiac arrest discharge locations. *Resuscitation plus*, 8:100185, 2021.
- Vincent Jeanselme, Glen Martin, Niels Peek, Matthew Sperrin, Brian Tom, and Jessica Barrett. Deepjoint: Robust survival modelling under clinical presence shift. *arXiv preprint arXiv:2205.13481*, 2022.
- Alistair EW Johnson, Tom J Pollard, Lu Shen, H Lehman Li-Wei, Mengling Feng, Mohammad Ghassemi, Benjamin Moody, Peter Szolovits, Leo Anthony Celi, and Roger G Mark. Mimic-iii, a freely accessible critical care database. *Scientific data*, 3(1):1–9, 2016.
- Luke K Kim, Patrick Looser, Rajesh V Swaminathan, James Horowitz, Oren Friedman, Ji Hae Shin, Robert M

- Minutello, Geoffrey Bergman, Harsimran Singh, S Chiu Wong, et al. Sex-based disparities in incidence, treatment, and outcomes of cardiac arrest in the united states, 2003–2012. *Journal of the American Heart Association*, 5(6):e003704, 2016.
- Agostina J Larrazabal, Nicolás Nieto, Victoria Peterson, Diego H Milone, and Enzo Ferrante. Gender imbalance in medical imaging datasets produces biased classifiers for computer-aided diagnosis. *Proceedings of the National Academy of Sciences*, 117(23):12592–12594, 2020.
- Yu-Kai Lin, Mingfeng Lin, and Hsinchun Chen. Do electronic health records affect quality of care? evidence from the hitech act. *Information Systems Research*, 30(1):306–318, 2019.
- Zachary C Lipton, David Kale, and Randall Wetzell. Directly modeling missing data in sequences with rnns: Improved classification of clinical time series. In *Machine Learning for Healthcare Conference*, pages 253–270, 2016.
- Roderick JA Little and Donald B Rubin. *Statistical analysis with missing data*, volume 793. John Wiley & Sons, 2019.
- Fernando Martínez-Plumed, Cèsar Ferri, David Nieves, and José Hernández-Orallo. Fairness and missing values. *arXiv preprint arXiv:1905.12728*, 2019.
- Karthika Mohan and Judea Pearl. Graphical models for processing missing data. *Journal of the American Statistical Association*, 116(534):1023–1037, 2021.
- Chirag Nagpal, Vincent Jeanselme, and Artur Dubrawski. Deep parametric time-to-event regression with time-varying covariates. In Russell Greiner, Neeraj Kumar, Thomas Alexander Gerds, and Mihaela van der Schaar, editors, *Proceedings of AAAI Spring Symposium on Survival Prediction - Algorithms, Challenges, and Applications 2021*, volume 146 of *Proceedings of Machine Learning Research*, pages 184–193. PMLR, 22–24 Mar 2021. URL <http://proceedings.mlr.press/v146/nagpal21a.html>.
- Alan Nelson. Unequal treatment: confronting racial and ethnic disparities in health care. *Journal of the national medical association*, 94(8):666, 2002.
- Craig D Newgard and Jason S Haukoos. Advanced statistics: missing data in clinical research—part 2: multiple imputation. *Academic Emergency Medicine*, 14(7):669–678, 2007.
- Craig D Newgard and Roger J Lewis. Missing data: how to best account for what is not known. *Jama*, 314(9):940–941, 2015.
- Todd G. Nick and Kathleen M. Campbell. *Logistic Regression*, pages 273–301. Humana Press, Totowa, NJ, 2007.
- Alejandro Noriega-Campero, Michiel A Bakker, Bernardo Garcia-Bulle, and Alex ‘Sandy’ Pentland. Active fairness in algorithmic decision making. In *Proceedings of the 2019 AAAI/ACM Conference on AI, Ethics, and Society*, pages 77–83, 2019.
- Keith Norris and Allen R Nissenson. Race, gender, and socioeconomic disparities in ckd in the united states. *Journal of the American Society of Nephrology*, 19(7):1261–1270, 2008.
- Stephen Pfohl, Ben Marafino, Adrien Coulet, Fatima Rodriguez, Latha Palaniappan, and Nigam H Shah. Creating fair models of atherosclerotic cardiovascular disease risk. In *Proceedings of the 2019*

- AAAI/ACM Conference on AI, Ethics, and Society*, pages 271–278, 2019.
- Eboni G Price-Haywood, Jeffrey Burton, Daniel Fort, and Leonardo Seoane. Hospitalization and mortality among black patients and white patients with covid-19. *New England Journal of Medicine*, 382(26):2534–2543, 2020.
- Alvin Rajkomar, Michaela Hardt, Michael D Howell, Greg Corrado, and Marshall H Chin. Ensuring fairness in machine learning to advance health equity. *Annals of internal medicine*, 169(12):866–872, 2018.
- Eliane Rööfli, Selen Bozkurt, and Tina Hernandez-Boussard. Peeking into a black box, the fairness and generalizability of a mimic-iii benchmarking model. *Scientific Data*, 9(1):1–13, 2022.
- Donald B Rubin. Inference and missing data. *Biometrika*, 63(3):581–592, 1976.
- Donald B Rubin. *Multiple imputation for nonresponse in surveys*, volume 81. John Wiley & Sons, 2004.
- Alexander Rusanov, Nicole G Weiskopf, Shuang Wang, and Chunhua Weng. Hidden in plain sight: bias towards sick patients when sampling patients with sufficient electronic health record data for research. *BMC medical informatics and decision making*, 14(1):51, 2014.
- Maytal Saar-Tsechansky and Foster Provost. Handling missing values when applying classification models. *Journal of Machine Learning Research*, 2007.
- Laleh Seyyed-Kalantari, Guanxiong Liu, Matthew McDermott, Irene Y Chen, and Marzyeh Ghassemi. Chexclusion: Fairness gaps in deep chest x-ray classifiers. In *BIO-COMPUTING 2021: proceedings of the Pacific symposium*, pages 232–243. World Scientific, 2020.
- Anis Sharafoddini, Joel A Dubin, David M Maslove, and Joon Lee. A new insight into missing data in intensive care unit patient profiles: Observational study. *JMIR medical informatics*, 7(1):e11605, 2019.
- Sandip Sinharay, Hal S Stern, and Daniel Russell. The use of multiple imputation for the analysis of missing data. *Psychological methods*, 6(4):317, 2001.
- Rose Sisk, Lijing Lin, Matthew Sperrin, Jessica K Barrett, Brian Tom, Karla Diaz-Ordaz, Niels Peek, and Glen P Martin. Informative presence and observation in routine health data: A review of methodology for clinical risk prediction. *Journal of the American Medical Informatics Association*, 2020.
- Kyle T Smith, Denise Monti, Nageen Mir, Ellen Peters, Renuka Tipirneni, and Mary C Politi. Access is necessary but not sufficient: factors influencing delay and avoidance of health care services. *MDM Policy & Practice*, 3(1):2381468318760298, 2018.
- Matthew Sperrin, Glen P Martin, Rose Sisk, and Niels Peek. Missing data should be handled differently for prediction than for description or causal explanation. *Journal of Clinical Epidemiology*, 125:183–187, 2020.
- Ala Szczepura. Access to health care for ethnic minority populations. *Postgraduate medical journal*, 81(953):141–147, 2005.
- Evan J Tsiklidis, Talid Sinno, and Scott L Diamond. Predicting risk for trauma patients using static and dynamic information from the mimic iii database. *Plos one*, 17(1):e0262523, 2022.
- Birgit Vogel, Monica Acevedo, Yolande Appelman, C Noel Bairey Merz, Alaide Chieffo, Gemma A Figtree, Mayra Guerrero,

- Vijay Kunadian, Carolyn SP Lam, Angela HEM Maas, et al. The lancet women and cardiovascular disease commission: reducing the global burden by 2030. *The Lancet*, 397(10292):2385–2438, 2021.
- Shirly Wang, Matthew BA McDermott, Geeticka Chauhan, Marzyeh Ghassemi, Michael C Hughes, and Tristan Naumann. Mimic-extract: A data extraction, pre-processing, and representation pipeline for mimic-iii. In *Proceedings of the ACM Conference on Health, Inference, and Learning*, pages 222–235, 2020.
- Nicole G Weiskopf, Alex Rusanov, and Chunhua Weng. Sick patients have more data: the non-random completeness of electronic health records. In *AMIA Annual Symposium Proceedings*, volume 2013, page 1472. American Medical Informatics Association, 2013.
- Joel S Weissman, Robert Stern, Stephen L Fielding, and Arnold M Epstein. Delayed access to health care: risk factors, reasons, and consequences, 1991.
- Brian J Wells, Kevin M Chagin, Amy S Nowacki, and Michael W Kattan. Strategies for handling missing data in electronic health record derived data. *Egems*, 1(3), 2013.
- Ian R White, Patrick Royston, and Angela M Wood. Multiple imputation using chained equations: issues and guidance for practice. *Statistics in medicine*, 30(4):377–399, 2011.
- Angela M Wood, Ian R White, and Simon G Thompson. Are missing outcome data adequately handled? a review of published randomized controlled trials in major medical journals. *Clinical trials*, 1(4):368–376, 2004.
- Jinghong Xu, Li Tong, Jiyou Yao, Zilu Guo, Ka Yin Lui, XiaoGuang Hu, Lu Cao, Yanping Zhu, Fa Huang, Xiangdong Guan, et al. Association of sex with clinical outcome in critically ill sepsis patients: a retrospective analysis of the large clinical database mimic-iii. *Shock (Augusta, Ga.)*, 52(2):146, 2019.
- Ruqaiijah Yearby. Racial disparities in health status and access to healthcare: the continuation of inequality in the united states due to structural racism. *American Journal of Economics and Sociology*, 77(3-4):1113–1152, 2018.
- Haoran Zhang, Amy X Lu, Mohamed Abdalla, Matthew McDermott, and Marzyeh Ghassemi. Hurtful words: quantifying biases in clinical contextual word embeddings. In *proceedings of the ACM Conference on Health, Inference, and Learning*, pages 110–120, 2020.
- Haoran Zhang, Natalie Dullerud, Karsten Roth, Lauren Oakden-Rayner, Stephen Pfohl, and Marzyeh Ghassemi. Improving the fairness of chest x-ray classifiers. In Gerardo Flores, George H Chen, Tom Pollard, Joyce C Ho, and Tristan Naumann, editors, *Proceedings of the Conference on Health, Inference, and Learning*, volume 174 of *Proceedings of Machine Learning Research*, pages 204–233. PMLR, 07–08 Apr 2022.
- Yiliang Zhang and Qi Long. Fairness in missing data imputation. *arXiv preprint arXiv:2110.12002*, 2021.
- Xiao-Hua Zhou, George J Eckert, and William M Tierney. Multiple imputation in public health research. *Statistics in medicine*, 20(9-10):1541–1549, 2001.

Appendix A. Experiments

This section provides additional details on the experimental design.

A.1. Simulation study

Data Generation. The proposed synthetic population consists of 10,000 points for the majority group and 100 for the marginalised group resulting in a sample size of $N = 10,100$ with a ratio of 100:1. Each individual is represented in this dataset as a pair of covariates, i.e. $X \in \mathbb{R}^2$. For each group, 50% presents the condition, i.e. $\mathbb{P}(Y_i = 0) = 0.5$. Negatives are drawn from the normal distribution $\mathcal{N}((0, 0), 0.25)$. The disease characterisation, i.e. the boundary between positive and negatives, differs between groups with positive from the majority (resp. the marginalised group) sampled from $\mathcal{N}((1, 0), 0.25)$ (resp. $\mathcal{N}((0, 1), 0.25)$). Figure 6 shows the density distribution of the generated population.




Figure 6: Density distributions of the generated population.

Missingness. In this synthetic population, 50% of the dimension X_2 is removed in a given subgroup to enforce the three clinical presence scenarios. We enforce the following clinical missingness:

- Limited access to quality care (S1):

$$O_2 \mid [G = 1] \sim \text{Bern}(0.5)$$

- (Mis)-informed collection (S2):

$$O_2 \mid [X_1 > 0.5] \sim \text{Bern}(0.5)$$

- Confirmation bias (S3):

$$O_2 \mid [X_2 > 0.5] \sim \text{Bern}(0.5)$$

With O_2 the observation indicator associated with X_2 and G , the group membership ($G = 1$ indicates a member of the marginalised group).

Modelling. We generate 100 datasets and enforce the different missingness patterns before running a logistic regression with an l2 penalty ($\lambda = 1$). Results are computed on the 20% test set and averaged over the 100 iterations with 95% confidence bounds reported.

Tabular results. Table 3 presents the AUC-ROC for the minority and majority groups and the different imputation strategies.

Missing Indicators. Similarly, Table 4 presents the AUC-ROC when a missing indicator is added for modelling (as proposed in Section 4.2). These results echos similar points to the main paper:

- State-of-the-art methodologies can perform similarly at the population level but harm marginalised groups differently as shown in (S1).
- No methodology consistently outperforms the others.
- Recommendation of adding missing indicators can hurt performance as MICE and Group MICE show in (S2).
- All methodologies benefit in (S3) in which the missingness is informative of the missing value itself.

Table 3: AUC-ROC divided by scenarios, group and imputation strategy - Mean (std) over 100 simulations.

Scenario	Group	Imputation strategy		
		Median	MICE	Group MICE
(S1)	Majority	0.997 (0.000)	0.997 (0.000)	0.997 (0.000)
	Minority	0.672 (0.026)	0.629 (0.026)	0.733 (0.026)
	Population	0.995 (0.000)	0.995 (0.000)	0.995 (0.000)
(S2)	Majority	0.997 (0.000)	0.997 (0.000)	0.997 (0.000)
	Minority	0.785 (0.023)	0.795 (0.022)	0.758 (0.024)
	Population	0.995 (0.000)	0.995 (0.000)	0.994 (0.000)
(S3)	Majority	0.876 (0.002)	0.945 (0.001)	0.947 (0.001)
	Minority	0.557 (0.030)	0.587 (0.029)	0.577 (0.029)
	Population	0.873 (0.002)	0.942 (0.001)	0.943 (0.001)

Table 4: AUC-ROC divided by scenarios, group and imputation strategy with missing indicators - Mean (std) over 100 simulations.

Scenario	Group	Imputation strategy with missing indicator		
		Median	MICE	Group MICE
(S1)	Majority	0.997 (0.000)	0.997 (0.000)	0.997 (0.000)
	Minority	0.684 (0.026)	0.641 (0.026)	0.661 (0.026)
	Population	0.995 (0.000)	0.995 (0.000)	0.994 (0.000)
(S2)	Majority	0.997 (0.000)	0.997 (0.000)	0.997 (0.000)
	Minority	0.798 (0.022)	0.797 (0.021)	0.764 (0.024)
	Population	0.995 (0.000)	0.995 (0.000)	0.994 (0.000)
(S3)	Majority	0.993 (0.000)	0.992 (0.000)	0.992 (0.000)
	Minority	0.694 (0.026)	0.700 (0.026)	0.698 (0.026)
	Population	0.990 (0.000)	0.989 (0.001)	0.989 (0.000)

A.2. MIMIC III

Dataset. After preprocessing (Wang et al., 2020) and standardisation, the MIMIC III dataset consists of 36,296 patients with 67 different laboratory tests. Focusing on the three marginalised groups of interest, the population can be further divided into marginalised subgroups as presented in Figure 7. First, this real-world distribution justifies our simulation choice as group intersection might be underrepresented in a dataset and present different characteristics than the majority. Second, this underlines the problem of identifying marginalised groups as they can be impacted differently by the same model.




Figure 7: Venn diagram of the population distribution in the three marginalised groups.

Further clinical presence evidence.

Table 1 shows the difference in testing at the end of the 24-hour observation period, one can study the temporal evolution of the testing procedure. Clinical presence is expressed not only in the missingness but in the temporality of the generative process (Jeanselme et al., 2022). Figure 8 displays the temporal

evolution of the number of laboratory tests performed for both survivors and patients who die after the observation period. This motivates our future work on studying the impact of strategies handling irregular time series on algorithmic fairness.




Figure 8: Evolution of the number of tests during the observation period divided by outcome.

Experimental design. For this real-world dataset, patients are split into three groups: 80% for training, 10% for validation and 10% for hyper-parameters selection. The hyper-parameter search consisted of the l2 penalty selection for the logistic regression among $\lambda \in [0.1, 1., 10., 100.]$.

We bootstrapped the test set 100 times and report the mean and 95% confidence bounds.

Tabular results. Table 5 presents the AUC-ROC for each group and imputation strategy. Similarly Table 6 (resp. 7) shows the prioritisation (resp. the false negative rate).

Table 5: AUC-ROC performance divided by group and imputation strategy - Bootstrapped mean (std). Bold indicates the highest AUC-ROC.

Group	Imputation strategy		
	Median	MICE	Group MICE
Black	0.805 (0.030)	0.781 (0.036)	0.765 (0.037)
Non Black	0.744 (0.010)	0.740 (0.010)	0.738 (0.010)
Female	0.735 (0.016)	0.733 (0.018)	0.730 (0.018)
Male	0.757 (0.016)	0.748 (0.014)	0.744 (0.014)
Public	0.739 (0.011)	0.726 (0.012)	0.722 (0.013)
Private	0.751 (0.021)	0.771 (0.018)	0.759 (0.021)
Population	0.748 (0.011)	0.742 (0.012)	0.738 (0.012)

Table 6: Prioritisation rate divided by group and imputation strategy - Bootstrapped mean (std). Bold indicates the highest prioritisation rate.

Group	Imputation strategy		
	Median	MICE	Group MICE
Black	0.343 (0.017)	0.323 (0.017)	0.150 (0.013)
Non Black	0.296 (0.004)	0.298 (0.004)	0.313 (0.005)
Female	0.328 (0.008)	0.320 (0.007)	0.315 (0.007)
Male	0.280 (0.006)	0.285 (0.006)	0.289 (0.006)
Public	0.350 (0.006)	0.342 (0.006)	0.354 (0.007)
Private	0.200 (0.007)	0.215 (0.007)	0.193 (0.006)
Population	0.299 (0.004)	0.299 (0.004)	0.300 (0.004)

Table 7: False Negative rate divided by group and imputation strategy - Bootstrapped mean (std). Bold indicates the smallest FNR.

Group	Imputation strategy		
	Median	MICE	Group MICE
Black	0.298 (0.076)	0.314 (0.078)	0.520 (0.081)
Non Black	0.365 (0.016)	0.370 (0.018)	0.356 (0.019)
Female	0.376 (0.030)	0.378 (0.030)	0.371 (0.029)
Male	0.346 (0.027)	0.357 (0.025)	0.367 (0.025)
Public	0.339 (0.021)	0.351 (0.018)	0.341 (0.020)
Private	0.425 (0.039)	0.412 (0.039)	0.458 (0.042)
Population	0.361 (0.019)	0.366 (0.020)	0.367 (0.020)

Table 8: AUC-ROC performance divided by group and imputation strategy - Bootstrapped mean (std). Bold indicates the highest AUC-ROC.

Group	Imputation strategy with missing indicators		
	Median	MICE	Group MICE
Black	0.827 (0.027)	0.817 (0.029)	0.814 (0.029)
Non Black	0.786 (0.010)	0.785 (0.010)	0.783 (0.010)
Female	0.770 (0.013)	0.772 (0.014)	0.769 (0.014)
Male	0.801 (0.013)	0.798 (0.013)	0.796 (0.013)
Public	0.773 (0.010)	0.770 (0.010)	0.767 (0.011)
Private	0.816 (0.016)	0.824 (0.016)	0.821 (0.017)
Population	0.789 (0.008)	0.788 (0.009)	0.786 (0.009)

Table 9: Prioritisation rate divided by group and imputation strategy - Bootstrapped mean (std). Bold indicates the highest prioritisation rate.

Group	Imputation strategy with missing indicators		
	Median	MICE	Group MICE
Black	0.317 (0.016)	0.318 (0.016)	0.194 (0.014)
Non Black	0.298 (0.005)	0.298 (0.005)	0.309 (0.005)
Female	0.320 (0.008)	0.314 (0.007)	0.315 (0.007)
Male	0.285 (0.006)	0.290 (0.006)	0.289 (0.006)
Public	0.341 (0.005)	0.338 (0.006)	0.339 (0.006)
Private	0.220 (0.006)	0.225 (0.005)	0.222 (0.005)
Population	0.300 (0.005)	0.300 (0.004)	0.300 (0.004)

Table 10: False Negative rate divided by group and imputation strategy - Bootstrapped mean (std). Bold indicates the smallest FNR.

Group	Imputation strategy with missing indicators		
	Median	MICE	Group MICE
Black	0.224 (0.063)	0.251 (0.065)	0.419 (0.071)
Non Black	0.300 (0.018)	0.287 (0.018)	0.285 (0.018)
Female	0.298 (0.023)	0.294 (0.026)	0.311 (0.027)
Male	0.296 (0.025)	0.272 (0.025)	0.277 (0.027)
Public	0.292 (0.020)	0.279 (0.018)	0.288 (0.019)
Private	0.302 (0.038)	0.301 (0.039)	0.321 (0.040)
Population	0.294 (0.017)	0.283 (0.018)	0.293 (0.018)

Missing Indicators. Similarly, Table 8 presents the AUC-ROC for each group and imputation strategy when missing indicators are added to the regression model. Table 9 (resp. 10) shows the prioritisation (resp. the false negative rate) for the same pipelines. This analysis shows that the missingness patterns are informative of the outcome of interest as adding the missing indicators as regressors improves performance. Note, however, that MICE Group presents lower AUC performance than MICE in Table 8. This observation echos the criticism of adding covariate for a more plausible MAR assumption. Finally, these tables underline how no method consistently outperforms the others across groups and metrics.

Threshold sensitivity. In Section 5.2, we present results for a policy of 30% additional care. As we arbitrarily chose this threshold, we propose to measure how the results vary under two different thresholds: 5% and 50%. Figures 9, 10 and 11 present the results at 5%, 30% and 50% thresholds. First, note that the magnitude of the Δ in prioritisation increases with larger thresholds, but similar trends are observed. This indicates that members of the same group have similar risk scores. Increasing the threshold, therefore, further penalises this whole group. Second, the Δ in false positive rates demonstrates how the choice of imputation is sensitive to the target task. In addition to validating the insights from Section 5.2, this set of experiments demonstrates that the target task may also affect whether an imputation methodology favours or penalises a given group.




Figure 9: Prioritisation performance gaps Δ across marginalised groups in MIMIC III experiment for 5% additional care. If $\Delta > 0$, the marginalised group has a larger value of the given metric than the rest of the population.




Figure 10: Prioritisation performance gaps Δ across marginalised groups in MIMIC III experiment for 30% additional care. If $\Delta > 0$, the marginalised group has a larger value of the given metric than the rest of the population.




Figure 11: Prioritisation performance gaps Δ across marginalised groups in MIMIC III experiment for 50% additional care. If $\Delta > 0$, the marginalised group has a larger value of the given metric than the rest of the population.