Sharing Time-to-Event Data with Privacy Protection

Luca Bonomi
Dept. of Biomedical Informatics
Vanderbilt University
Nashville, TN
luca.bonomi@vumc.org

Liyue Fan
Dept. of Computer Science
University of North Carolina at Charlotte
Charlotte, NC
liyue.fan@uncc.edu

***Abstract*—Sharing time-to-event data is beneficial for enabling collaborative research efforts (e.g., survival studies), facilitating the design of effective interventions, and advancing patient care (e.g., early diagnosis). Despite numerous privacy solutions for sharing time-to-event data, recent research studies have shown that external information may become available (e.g., self-disclosure of study participation on social media) to an adversary, posing new privacy concerns. In this work, we formulate a cohort inference attack for time-to-event data sharing, in which an informed adversary aims at inferring the membership of a target individual in a specific cohort. Our study investigates the privacy risks associated with time-to-event data and evaluates the empirical privacy protection offered by popular privacy-protecting solutions (e.g., binning, differential privacy). Furthermore, we propose a novel approach to privately release individual level time-to-event data with high utility, while providing indistinguishability guarantees for the input value. Our method TE-Sanitizer is shown to provide effective mitigation against the inference attacks and high usefulness in survival analysis. The results and discussion provide domain experts with insights on the privacy and the usefulness of the studied methods.**

***Keywords*-Provable Privacy, Time-to-event Data, Survival Analysis**

I. Introduction

Temporal health data have been increasingly collected and shared to promote data-driven research and to advance patient care. For example, the temporal data about the COVID-19 pandemic [1] have made it possible for the research community to gain useful insights about the infection rate of the virus, as well as its impact on vulnerable populations [2]–[4]. Among all applications, the sharing of time-to-event data can provide great benefits in population and epidemiology studies [5]–[7]. Specifically, time-to-event data capture the length of time until the occurrence of certain clinical events (e.g., survival study events and hospital discharges), which can be utilized to estimate a survival function and to facilitate the design of effective interventions. As an example, recent research conducted survival analysis on time-to-event data collected during the COVID-19 pandemic to predict patient length of stay [8] as well as time until severe outcomes [9].

In time-to-event analysis (e.g., survival studies), time-to-event data are shared at aggregate level (e.g., survival functions) or at individual level in an de-identified manner, and therefore do not contain PHI information about individuals. However, an adversary may leverage publicly available data or background knowledge to infer sensitive information about target individuals in the data. Our prior work [10] shows that a powerful adversary may infer an individual's participation in the data by observing updates in the survival curve. While such an attack may be challenging to carry out, the privacy risk exists inherently. In fact, recent studies have shown that data contributors may inadvertently disclose information about their participation to research studies on social media, such as the study name and the time of data contribution [11], [12]. An informed adversary may utilize the acquired information about the target individual as well as the shared study dataset and learn the target's membership in a specific cohort (e.g., case group), thus revealing the target's phenotype data.

To address these privacy concerns, several privacy-protecting solutions have been proposed for sharing time-to-event data in biomedical applications [10], [13]–[15]. Among them, binning with suppression (also known as thresholding) is commonly used in applications. In this approach, time-to-events are aggregated over disjoint temporal bins and bin counts are reported only if they meet a predetermined threshold (e.g., more than 5 patients per bin). While this approach is popular and intuitive, it may lead to the loss of time-to-events where the threshold is not met. To mitigate the data loss, researchers have proposed perturbation-based approaches, in which the output results are perturbed with random noise [14], [15]. However, these approaches do not provide provable privacy guarantees. A recent approach [10] adopted differential privacy [16] in privacy-protecting survival studies. In this privacy model, the output results are perturbed with calibrated random noise to ensure that they are "indistinguishable" regardless the presence or absence of any individual data contributor. Due to the strong privacy guarantee, the injected random noise may limit the clinical usefulness of the shared data.

The goal of this work is two-fold. First, we propose a new privacy-protecting solution for sharing time-to-event data at the individual level. Our solution achieves time-to-event indistinguishability (TI), a relaxed privacy notion compared

to standard differential privacy, in which indistinguishability is guaranteed within a bounded temporal window. Evaluations demonstrate that our approach significantly improves the usability of the shared data compared to existing privacy methods, while providing provable privacy guarantees. Second, we investigate the privacy risk of sharing time-to-event data by formulating an inference attack, in which an informed adversary aims at inferring the cohort membership of a target individual in the study. Under this adversarial model, our evaluations show that existing privacy approaches may provide unbalanced privacy protection between cohorts, where strengthening the privacy for some cohorts may decrease the protection for others, and may provide limited usability when input dataset is small. Overall, our results with two real-world datasets provide domain experts (e.g., clinicians) with useful insights on the effectiveness and usability of current privacy solutions for time-to-event data, which may help them identify suitable privacy solutions for specific applications.

II. Methods

We consider a dataset of n individuals $D = \{(c_i, t_i)\}_{i=1}^{n}$, where each individual i is associated with a cohort label $c_i \in \{1, \ldots, C\}$, and a temporal value t_i denoting the time-to-event information, e.g., the length of time until the occurrence of a survival study event or hospital discharge. While other covariates may be recorded for each individual, they are omitted in the problem definition for brevity and we will discuss the privacy implications of additional covariates in Section IV.

Beyond representing the study cohort, c_i may represent the result of patient grouping either via additional analysis (e.g., cluster analysis using genomic data [17]) or by attribute values (e.g., age group, gender). While privacy methods could be applied to hide individual cohort membership (e.g., perturbing the cohort label), preserving the truthfulness of cohort information is essential for applications. Therefore, in this work, we investigate the privacy risks and solutions for sharing individual-level time-to-event t_i, while preserving the cohort label c_i to maximize data truthfulness.

A. Current Privacy Solutions

Current solutions for time-to-event data sharing build on a variety of privacy techniques [14], [18]–[20]. Below, we summarize the approaches most relevant to the problem studied in this manuscript.

Binning-based Privacy Methods. A common privacy practice for sharing time-to-event data relies on aggregation and statistical disclosure control, where an individual is hidden within a sufficiently large group. Specifically, this privacy strategy produces aggregate-level statistics (e.g., number of individuals in cohort c at time-to-event t), which are shared with external users (e.g., researchers). As an example, binning with suppression produces aggregate counts of time-to-event data over disjoint temporal bins of fixed length (e.g., temporal interval of 10 time units), where only bin counts larger than a threshold are released (e.g., at least 5 individuals in the bin). Intuitively, in binning with suppression, the number of time-to-events is shared only if a sufficient number of time-to-events occur in the same bin.

Perturbation-based Privacy Methods. Over the years, researchers have proposed a variety of methods that achieve privacy via data perturbation. As an example, studies have developed methods to consistently shift the dates to hide the original temporal information in each record. This privacy technique hides the absolute date values and preserves relative information between records (e.g., the gap between two visits). However, it is challenging the quantify the privacy protection provided by shifting. Among the perturbation-based approaches, the differential privacy [16] (DP) model has shown to provide rigorous privacy protection for aggregate-level data sharing. In this model, privacy is achieved by injecting random noise in the output results. The magnitude of the noise is carefully calibrated to hide the presence of any individual in the data. Our recent study [10] has demonstrated the applicability of the DP model in sharing aggregate time-to-event data for survival studies.

Randomized Response Privacy Methods. To enable fine-grained data sharing, researchers have proposed new privacy methods that extend the traditional differential privacy model to individual-level data. An example is the local differential privacy model (LDP), in which individual-level data are protected using randomized response methods (e.g., in structured survey interviews [21]). Recent privacy works have shown that LDP methods can be used to build strong privacy solutions and provide individual-level indistinguishability for any input data value [22]–[24]. Despite promising results, it is challenging to preserve individual-level data usability when the input domain is large. For example, the staircase mechanism [25] may lead to output results with high variance for strong privacy, resulting in output data values that are significantly different from the original data, thus reducing the usefulness of the shared data. To address this limitation, the metric privacy model [26] extends the standard differential privacy model over generic metric spaces, enabling the privacy mechanism to provide a more flexible privacy protection. While this protection may be weaker than traditional DP, the shared data tend to preserve better usability. As an example, when applied to the geospatial domain, the metric privacy model can significantly improve the usability of shared locations [27], while providing a reasonable privacy protection.

B. Time-to-event Sanitization with TE-Sanitizer

In this section, we present our approach for time-to-event data sharing. There are two key differences that

distinguish our approach from existing privacy solutions. (1) Our method achieves privacy at individual-level by perturbing the recorded time-to-event t_i, while preserving the clinical information in the data (i.e., cohort label c_i) for applications. (2) Our method builds on the principle of randomized response to protect the individual time-to-event data. Compared to standard LDP, our approach constrains the output space to a temporal window, which helps retain the usefulness of the data.

We first define the indistinguishability notion for time-to-event data. Intuitively, given a window size W and any sanitized time-to-event $\hat{t}$, the privacy method should provide indistinguishability for any pair of t_i, t_j in input that are within W time units from $\hat{t}$. Formally,

Definition II.1 (Time-to-event Indistinguishability (TI))**.** For any output $\hat{t}$ and any pair of time-to-events $t_i, t_j \in D$ such that $|t_i - \hat{t}| \leq W$ and $|t_j - \hat{t}| \leq W$, a mechanism M satisfies (ϵW)-TI, if and only if the following inequality holds:

$$\frac{Pr[M(t_i) = \hat{t}]}{Pr[M(t_j) = \hat{t}]} \leq e^{\epsilon W} \tag{1}$$

where W is a pre-defined window size in time units (e.g., days or months).

With the TI definition, time-to-events that are temporally close will be indistinguishable to an adversary who observes the output results of M, while time-to-events that are far apart may not be. In practice, users may choose a larger W (e.g., entire duration of the study) to provide stronger privacy, i.e., indistinguishability for all time-to-event pairs. We vary W values in our evaluation to study the effects of the parameter empirically.

Our overall solution is outlined in Algorithm 1. Given an input time-to-event t and the window size W, our method randomly samples a sanitized time-to-event $\hat{t}$ according to the following probabilities:

$$Pr[M(t) = \hat{t}] = \begin{cases} 0 & \text{if } |t - \hat{t}| > W \\ e^{-\epsilon(t-\hat{t})}/(1+e^{-\epsilon}) & \text{if } t - \hat{t} = W \\ e^{-\epsilon(\hat{t}-t)}/(1+e^{-\epsilon}) & \text{if } \hat{t} - t = W \\ \frac{1-e^{-\epsilon}}{1+e^{-\epsilon}} e^{-\epsilon|t-\hat{t}|} & \text{otherwise} \end{cases} \tag{2}$$

The sampling method can be seen as an adaptation of the Truncated Symmetric Geometric Mechanism, which has shown to achieve near-optimal utility [28]. The sanitized dataset $\hat{D}$ in Algorithm 1 contains records with the obfuscated time-to-event $\hat{t}$ instead of t, with the cohort label unchanged.

The following result states the privacy guarantee provided by our solution.

Theorem II.1. Algorithm 1 satisfies (ϵW)-TI.

Algorithm 1 TE-Sanitizer

```
procedure TE-SANITIZER(D,ε,W)                    ▷
  Sanitize the input dataset D, with privacy parameter ε,
  and time window W
    D̂ ← ∅
    for (t, c) in D do
        Sample t̂ using equation (2)
        D̂ = D̂ ∪ {(t̂, c)}
    end for
    return D̂                          ▷ The sanitized dataset
end procedure
```

Proof: (Sketch.) Our algorithm uses the sampling procedure in equation (2) to generate the sanitized time-to-event as output. For any $\hat{t}$ and t_i, t_j such that $|t_i - \hat{t}| \leq W$ and $|t_j - \hat{t}| \leq W$, we analyze four possible cases below. The remaining cases are symmetric thus omitted.

- Case (a): $t_i > t_j > \hat{t}$, $0 < t_i - \hat{t} < W$, and $0 < t_j - \hat{t} < W$

$$\frac{Pr[M(t_i) = \hat{t}]}{Pr[M(t_j) = \hat{t}]} = \frac{e^{-\epsilon(t_i - \hat{t})}}{e^{-\epsilon(t_j - \hat{t})}} = \\ = e^{-\epsilon(t_i - t_j)} \leq e^{\epsilon W}$$

- Case (b): $t_i > t_j$, $0 < t_i - \hat{t} < W$, and $0 < \hat{t} - t_j < W$

$$\frac{Pr[M(t_i) = \hat{t}]}{Pr[M(t_j) = \hat{t}]} = \frac{e^{-\epsilon(t_i - \hat{t})}}{e^{-\epsilon(\hat{t} - t_j)}} = \\ = e^{-\epsilon(t_i - 2\hat{t} + t_j)} \leq e^{\epsilon W}$$

- Case (c): $t_i > t_j$, $t_i - \hat{t} = W$, and $\hat{t} - t_j = W$

$$\frac{Pr[M(t_i) = \hat{t}}{Pr[M(t_j) = \hat{t}]} = \frac{e^{-\epsilon(t_i - \hat{t})}}{e^{-\epsilon(\hat{t} - t_j)}} = \\ = e^{-\epsilon(t_i - 2\hat{t} + t_j)} = 1 \leq e^{\epsilon W}$$

- Case (d): $t_i > t_j$, $0 < t_i - \hat{t} < W$, and $\hat{t} - t_j = W$

$$\frac{Pr[M(t_i) = \hat{t}]}{Pr[M(t_j) = \hat{t}]} = \frac{(1 - e^{-\epsilon})e^{-\epsilon(t_i - \hat{t})}}{e^{-\epsilon(\hat{t} - t_j)}} = \\ (1 - e^{-\epsilon})e^{-\epsilon(t_i - \hat{t} - W)} \leq e^{\epsilon W}$$

This concludes the proof. ■

Theorem II.1 shows that our method provides quantifiable privacy protection for individual-level time-to-event data. Moreover, the user-defined window size W bounds the noise introduced in the data, thus improving the usability of the shared data. We will demonstrate the benefits of our proposed approach with extensive empirical evaluations in Section III.

C. Measuring Privacy Risk

In this work, we assume an informed adversary who: (1) has access to the sanitized dataset $\hat{D}$, (2) has prior knowledge about a target individual's participation to a study (i.e., in dataset D) and their time-to-event t, and (3) has knowledge of the privacy mechanism M used to generate $\hat{D}$ (i.e., how privacy is achieved, values of privacy parameters etc.) [29]. Our work protects the privacy of data contributors against untrusted data recipients, hence the first assumption. The second assumption is motivated by recent studies which show that individuals may accidentally disclose their participation to research studies, for example, in social media [11]. Therefore, the risk of an adversary possessing prior information about the target inherently exists. The third assumption has been suggested by Kifer and Lin [30], in which all information about the privacy mechanism M should be assumed public (i.e., privacy should not be achieved by obscurity).

Under those assumptions, the adversary's goal is to infer the cohort c of the target individual, upon receiving the sanitized data $\hat{D}$ with obscured time-to-event values. Specifically, given a target individual's time-to-event $t \in D$ and the dataset $\hat{D} = M(D)$ that has been sanitized by a privacy mechanism M, an adversary aims at inferring the target's cohort c by linking the real time-to-event t with published time-to-events $\hat{t}$ in $\hat{D}$.

Inference Attack. In our setting, we consider the risk of inferring the cohort c of a target individual with time-to-event t. Specifically, the adversary may compute a likelihood score for reconstructing the cohort with $\hat{c} \in \{1, ..., C\}$ as follows:

$$CL(\hat{c}, t) = \sum_{\hat{t}} Pr[\hat{c}|\hat{t}]Pr[\hat{t}|t] \tag{3}$$

where $Pr[\hat{c}|\hat{t}]$ can be estimated by observing the occurrences of cohort $\hat{c}$ among records with time-to-event $\hat{t}$ in $\hat{D}$, and $Pr[\hat{t}|t]$ is equivalent to Equation 2 (i.e., knowledge of the mechanism). In practice, the attacker may assign the target to the most likely cohort, i.e., $\hat{c}^* = \arg\max_{\hat{c}=1,...,C} \{CL(\hat{c}, t)\}$.

III. Evaluations

This section describes our evaluation methodology and results. Note that additional details and results can be found in the Supplementary Material which is included in the full version of the manuscript [31].

A. Methodology

Datasets. For empirical evaluation, we use two real-world datasets: METABRIC [32] and COVID-19 [33]. The METABRIC dataset contains 2,509 patients diagnosed with breast cancer. Among all the patients, we consider 1,444 patients who are divided into three different cohorts representing different cancer stages: cohort 1 (stage 1) with 501 patients, cohort 2 (stage 2) with 825 patients, and cohort 3 (stage 3) with 118 patients. The time-to-event in this dataset represents the patient's survival event duration in months. The COVID-19 dataset is an epidemiological dataset from the COVID-19 outbreak. We process the data as suggested in [8], resulting in a total of 186,396 patients divided in four cohorts according to age group: cohort 3 (age $>$ 60) with 42,786 patients, cohort 2 (47 $<$age$\leq$ 60) with 65,739 patients, cohort 1 (35 $<$age$\leq$ 46) with 17,711 patients, and cohort 0 (0 $\leq$age$\leq$ 35) with 60,160 patients. The time-to-event in this dataset represents the patient length of stay in days.

Comparisons. We evaluate our proposed technique (TE-Sanitizer) as well as other privacy solutions that are commonly used in practical applications: binning with suppression and differential privacy. Each method and its parameter(s) are described below. The adopted parameter values are reported in Table I.

- `TE-Sanitizer`(ϵ, W): this is our proposed privacy method, which sanitizes each record's time-to-event to achieve (ϵW)-TI.
- `BinSup`$(timeBin, sizeBin)$: this privacy method produces aggregate counts for each cohort about time-to-events over disjoint temporal intervals (i.e., bins of $timeBin$ time units), where the count for each bin is released if it is at least $sizeBin$ (i.e., count threshold) and 0 otherwise. In other words, it only reports counts for temporal intervals that contain a sufficient number of time-to-events or it returns a zero answer. The released aggregates are used to generate the sanitized data $\hat{D}$: sanitized records are created for bins with non-zero counts, each with the real cohort and the time-to-event rounded to the start of the interval.
- `DPTime`(ϵ): this approach builds on the original algorithm proposed in [10] to conduct survival analysis for non-parametric models (e.g., Kaplan-Meier). Here we adapt the algorithm such that it first computes aggregate statistics satisfying ϵ-differential privacy and then generates sanitized data $\hat{D}$ as done in BinSup.

Table I
Algorithm Parameters. Default values are boldfaced.

Parameter	Description	Values
ϵ	Privacy parameter for TE-Sanitizer and DPTime	$[0.1, 0.2, 0.4, \mathbf{0.8}, 1.6, 3.2]$
W	Size of the temporal window for TE-Sanitizer	$[5, \mathbf{10}, 25, 50, 100]$
TimeBin	Length of the temporal bins for BinSup	$[1, 5, \mathbf{10}]$
SizeBin	Count threshold of each temporal bin for BinSup	$[1, \mathbf{5}, 10, 20]$

Utility Measure. To measure the utility of the sanitized data, we consider two metrics, namely the Kullback–Leibler divergence (KL) and mean absolute error (MAE) between the original and sanitized data. The Kullback–Leibler divergence (KL) quantifies the distance between the distributions of time-to-event of the original and the sanitized data. The

mean absolute error (MAE) captures the average distortion of the time-to-event in time units. For both measures, smaller values indicate stronger similarity between the original and sanitized data, thus higher usability of the sanitized data. A formal description for these measures is reported in the Supplementary Material.

Privacy Risk Measure. To measure the risk of cohort inference, we adopt a common setting used in previous membership inference studies (e.g., [34]). In our study, we use the cohort likelihood scores to assign a target individual to a cohort and measure the precision of the attacker in inferring the real cohort. Specifically, we construct a balanced test set $D_{bal} \subset D$ with an equal number of individuals sampled from each cohort, i.e., 100 patients per cohort. Given a target individual in D_{bal} with time-to-event t_i, we compute the cohort likelihood score $CL(c, t_i)$ for every $c \in \{1, \ldots, C\}$ as in Equation 3. Then, for each cohort c, we rank all individuals in the test set by $CL(c, t_i)$ and those with score greater than 95% individuals are assigned to cohort c. Finally, we compute the precision on the attacker's success in inferring the real cohort for individuals in D_{bal} and report the results among 100 random samples of D_{bal}.

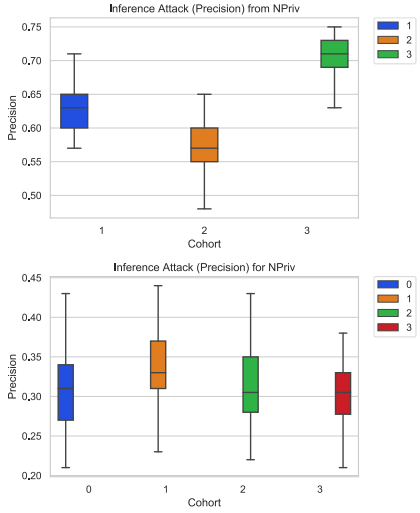

Figure 1. Adversary's precision of the inference attack in the non-private setting - NPriv. (Top) METABRIC dataset. (Bottom) COVID-19 dataset.

B. Privacy and Utility Evaluations

Here we study the privacy and utility performance for all privacy methods. We first report the privacy risk measures obtained in the non-private setting, i.e., releasing the real time-to-event for each individual. We then evaluate how method-specific privacy parameters impact the utility of the sanitized data and may mitigate the privacy risks compared to the non-private data.

1) Privacy risks in the non-private setting: Figure 1 reports the attacker's precision on the non-private data (i.e., when $\hat{D} = D$). The boxplots also describe the 95% confidence interval for each cohort. Note that for non-private data, the CL score can be simplified as $CL(\hat{c}, t) = Pr[\hat{c}|t]$, where $Pr[\hat{c}|t]$ is estimated from D for any $\hat{c}$. We observe that the attack is more successful in METABRIC. For example, the median precision for cohort 3 in METABRIC is above 70%, which is significantly higher than random guess (i.e., 33% for 3 cohorts). We believe that the larger size of the COVID-19 data helps reduce the privacy risk, as individuals from multiple cohorts may share the same time-to-event value. However, the privacy risk in COVID-19 is not negligible, as the median precision for every cohort is higher than random guess (i.e., 25% for 4 cohorts).

2) Evaluations on METABRIC Data: Figure 2 reports the privacy and utility results on the METABRIC data for all privacy methods. Note that additional results on MAE and parameter study are reported in Figure 7, Figure 8, and Figure 9 in Supplementary Material. For TE-Sanitizer, we observe that smaller values of the privacy parameter ϵ can greatly reduce the attack precision. As an example, for $\epsilon = 0.1$, the adversary's precision is reduced by roughly 15% for all cohorts compared to the non-private setting. We also observe that the sanitized data by TE-Sanitizer retain good utility. For example, the KL divergence between the original and sanitized data is quite low across all ϵ settings. Additionally, TE-Sanitizer leads to MAE$\leq$ 7 time units (Figure 9 in Supplementary Material), indicating high usability for individual-level data, whereas BinSup and DPTime inflict much higher MAE errors due to data loss and perturbation, respectively. Moreover, the parameter W has less impact on usability and privacy compared to the parameter ϵ (Figure 7 in Supplementary Material).

Regarding the DPTime method, which satisfies differential privacy, we observe lower attack precision compared to TE-Sanitizer (i.e., indicating stronger empirical privacy protection). However, this enhanced privacy protection comes at the expense of higher utility loss. In fact, we observe that the sanitized data produced by DPTime exhibit significantly higher KL divergence and MAE (Figure 9 in Supplementary Material) compared to the results of TE-Sanitizer.

For BinSup, we observe that the parameter $SizeBin$, which controls the threshold for the number of individuals in each bin, can have complex implications on the privacy protection in each cohort. For example, for cohort 1 and 3, the attack precision decreases significantly as $SizeBin$ increases. However, this parameter may have adverse effects on cohort 2, where larger values of $SizeBin$ improve the adversary's precision. Our evaluations suggest that this method may suppress more counts in cohorts 1 and 3 and preserve the data in cohort 2, as cohort 2 contains the majority of the data. Regarding the usability, KL divergence tends to increase with $SizeBin$, as more temporal events may be

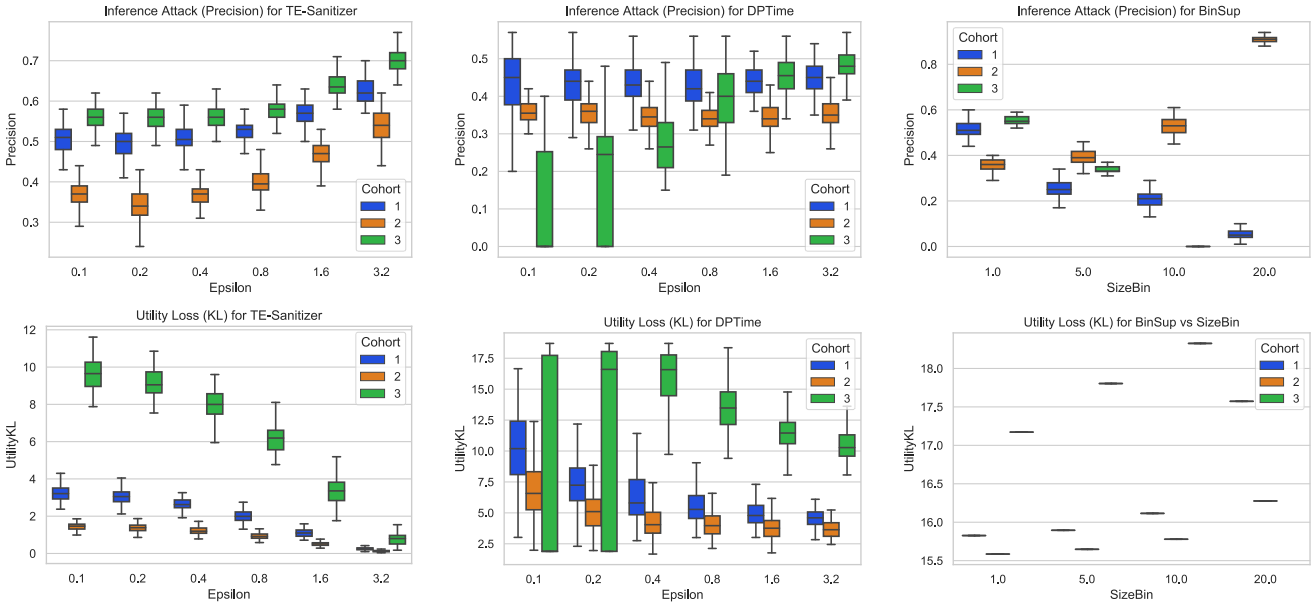

Figure 2. Evaluations on METABRIC - impact of the privacy parameters. (Top) Adversary's precision of inference attack using the sanitized data. (Bottom) Utility of the sanitized data in KL divergence.

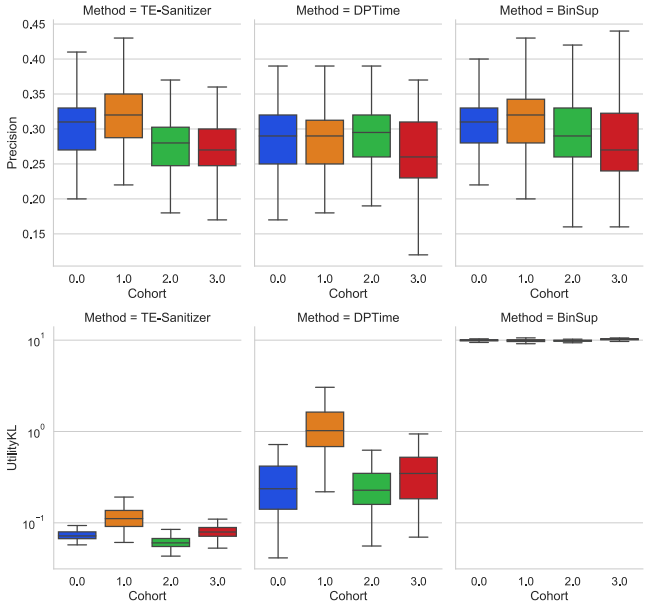

Figure 3. Evaluations on COVID-19 data. (Top) Adversary's precision of inference attack. (Bottom) Utility of the sanitized data in KL divergence.

suppressed [1]. The suppression of events also inflicts a large utility loss in terms of MAE (Figure 9 in Supplementary Material). Overall, the sanitized data exhibit a utility loss that is comparable to the standard differential privacy solution

[1]The KL divergence for cohort 3 at $SizeBin = 20$ is not reported, as all data for the cohort have been removed

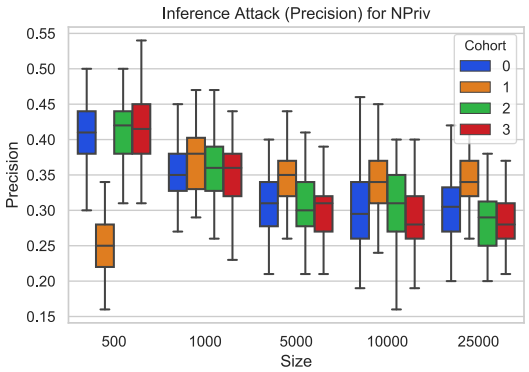

Figure 4. Precision of the inference attack for different size of the data measured on the original COVID-19 data (NP).

(i.e., DPTime). Additionally, we investigate the impact of the $TimeBin$ parameter on the privacy and usability of the sanitized data (Figure 8 in Supplementary Material). Our evaluations show that larger temporal intervals can reduce the attack precision but may inflict a significant utility loss for the sanitized data.

3) Evaluations on COVID-19 Data: For brevity, we present in Figure 3 a subset of results obtained with the default parameters and 10k individuals sampled from the COVID-19 dataset. The complete results with varying privacy parameters are reported in Figure 10 in Supplementary Material.

In Figure 3, all three privacy methods provide similar privacy protection in attack precision. Specifically, the me-

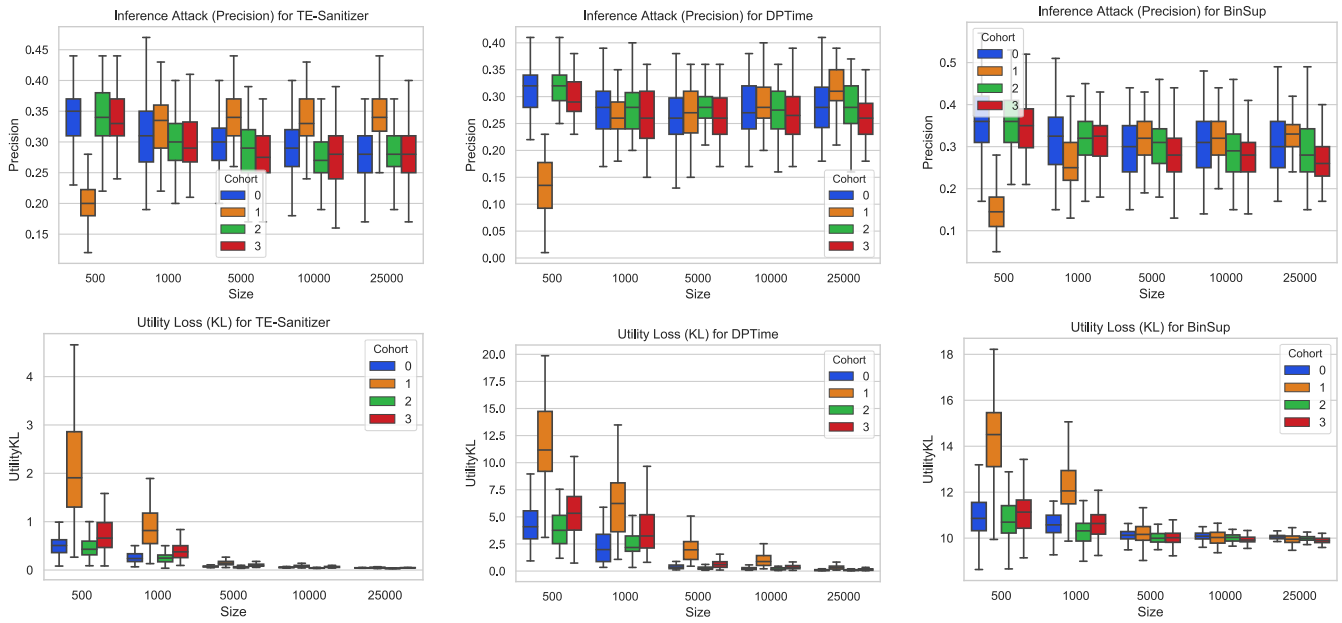

Figure 5. Evaluations with different number of individuals in COVID-19 dataset. (Top) Precision of the cohort inference attack with different size of the data, measured on the sanitized data. (Bottom) utility loss of sanitized data in KL divergence for different size of the data. For TE-Sanitizer and DPTime, we used $\epsilon = 1.0$. For BinSup, we set $BinSize = 5$ and $TimeBin = 2$.

dian attack precision ranges from 24% to 33%, whereas in the non-private setting the median precision is constantly higher than 30%. In terms of utility, we observe that TE-Sanitizer achieves the lowest KL divergence, demonstrating that the sanitized data are highly usable. Furthermore, DPTime method provides significantly better data usability compared to BinSup on this dataset. In fact, COVID-19 dataset contains a large number of individuals within a short time period, which suits well the differentially private DPTime method.

4) Input Data Size: We also study the impact of the input data size on the privacy and utility measures. In these evaluations, we vary the number of individuals in the range $[0.5k, 25k]$, randomly sampled without replacements from the COVID-19 dataset. For each dataset size, we report results over 100 sampled datasets of the given size.

We first report in Figure 4 the inference attack precision in the non-private setting. The attack precision in general decreases as the number of individuals in the dataset increases. In other words, the privacy risks are lessened in larger datasets. One exception we observe is that the attack precision for cohort 1 is low with 500 dataset size, i.e., 25% median precision. As the smallest cohort, very few patients from cohort 1 are sampled into a dataset of 500 patients; due to the short time period in COVID-19 data, those patients are likely to share time-to-event values with patients of other cohorts. As a result, $Pr[c =$'1'$|t]$ takes a small value for most t and patients of other cohorts have the same $Pr[c =$'1'$|t]$ score as cohort 1 patients when they have the same time-to-event value t, hence the low attack precision for cohort 1.

Figure 5 report the results with the sanitized data generated by privacy mechanisms. We observe that all privacy methods reduce the inference attack precision compared to the non-private setting, even when the dataset size is very small, i.e., 500. In fact, in order to achieve an attack precision $\leq$ 35% for every cohort, 500 dataset size is sufficient for TE-Sanitizer and DPTime, 1000 for BinSup; but 5000 dataset size is required in the non-private setting. Figure 5 also shows the utility loss decreases as more individuals are included in the data and TE-Sanitizer shows lower utility loss than other privacy methods, despite the size of input data. The utility results in MAE are reported in Figure 12 in Supplementary Material, in which we also observe a superior performance from TE-Sanitizer.

C. Case Study I: Survival Analysis for Breast Cancer Patients

In this case study, we evaluate the applicability of the privacy-protecting methods for time-to-event data in survival analysis for breast cancer patients. As a proof of concept, we adopt the Kaplan-Meier model, a non-parametric survival model that is extensively used in survival analysis [8], [35], [36], which estimates the survival probability directly from the time-to-event data. More details on the Kaplan-Meier model are provided in Supplementary Material.

Figure 6 top row depicts the comparison of survival curves obtained using different privacy methods on the

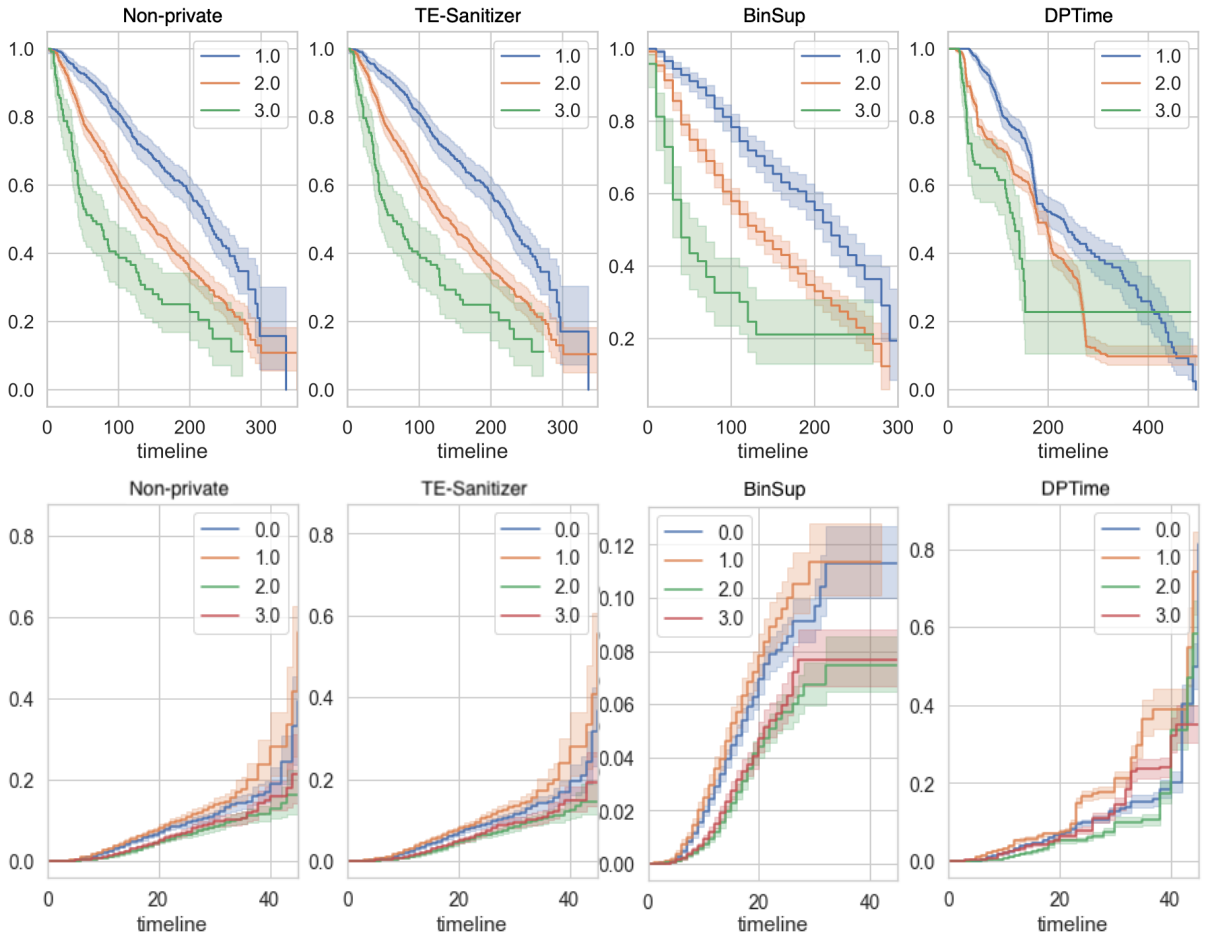

Figure 6. Case studies on the Kaplan-Meier analysis. (Top) Survival curves computed for the three cohorts $1, 2$, and 3 on non-private data, the sanitized data by TE-Sanitizer, BinSup, and DPTime. The parameters are: $\epsilon = 1.0$, $W = 10$, `Size_Bin=2`, and `Time_Bin=10`. (Bottom) Discharge curves (i.e., discharge-time probability) computed for the cohorts $0, 1, 2$, and 3 on the non-private data, the sanitized data by TE-Sanitizer, BinSup, and DPTime. The parameters are: $\epsilon = 1.0$, $W = 10$, `Size_Bin=5`, and `Time_Bin=1`.

METABRIC dataset. Overall, all the privacy methods are able to preserve the general trend of the survival function of each cohort. Among all the privacy methods, the survival curves computed on the sanitized data generated by TE-Sanitizer best resemble the non-private curves. Additionally, we measure similarity between the survival curves on the original and sanitized data using the log-rank test [37]. From the statistical results in Table II, only the curve for cohort 2 generated by DPTime solution is significantly different from the non-private curve.

Table II
LOG-RANK STATISTICS FOR THE KAPLAN-MEIER SURVIVAL CURVES WITH RESPECT TO THE NON-PRIVATE CURVES.

Methods	Cohort 1	Cohort 2	Cohort 3
TE-Sanitizer	0.0012	0.0013	0.0005
BinSup	0.2105	0.8558	1.5140
DPTime	1.9754	19.2220*	0.0074

*: significant at p-value≤ 0.05.

Table III
LOG-RANK STATISTICS FOR THE KAPLAN-MEIER DISCHARGE-TIME PROBABILITY CURVES WITH RESPECT TO THE NON-PRIVATE CURVES.

Methods	Cohort 0	Cohort 1	Cohort 2	Cohort 3
TE-Sanitizer	0.2429	0.0406	0.1380	0.0868
BinSup	3.4540	3.4539	4.7041*	5.3313*
DPTime	1.3207	31.6967*	1.6917	62.6090*

*: significant at p-value≤ 0.05.

D. Case Study II: Discharge Time for COVID-19 Patients

Similarly to the analysis in [8], we used the Kaplan-Meier model to estimate the discharge time for patients diagnosed with COVID-19 across different age groups. Figure 6 bottom row presents the results using the non-private data as well as the privacy-protecting methods considered in this paper. Our approach TE-Sanitizer produces discharge curves that best resemble those computed on the original data. In fact, the analysis results in Table III show no statistical difference

Table IV
OVERVIEW OF THE SANITIZATION METHODS FOR TIME-TO-EVENT DATA ACROSS EIGHT DIMENSIONS.

Method	Temporal Truthfulness	Cohort Size Truthfulness	Privacy Protection	Empirical Privacy Risk	Usability	Data Size	Input	Output
TE-Sanitizer	Time-to-events are perturbed	All individuals are preserved	Protect individual time-to-event value	Reduced w.r.t. the original data	Complex privacy model	Robust for varying data sizes	Time-to-event data	Time-to-event data
DPTime	Time-to-events are perturbed	The number of individuals is perturbed	Protect individual's presence	Reduced w.r.t. the original data	Complex privacy model	Large size is better	Time-to-event data	Aggregate Statistics
BinSup	Time-to-events are rounded	The number of individuals may be reduced	Hide individual by grouping and suppression	May not be reduced for all cohorts	Intuitive privacy notion	Large size is better	Time-to-event data	Aggregate Statistics

between the curves computed with our method and the non-private ones.

We can conclude that BinSup and DPTime have adverse effects on the usability of sanitized data, despite the large number of patients in the COVID-19 data (186,396 patients). Consider the BinSup method with $SizeBin = 5$ as illustrated in those results. Despite the small threshold, those suppressed bin counts inflict significant changes in the discharge curves for cohorts 2 and 3. For the DPTime method, the perturbation noise is equally distributed across all cohorts; as a result, cohorts with fewer individuals (i.e., cohort 1 and cohort 3) are more affected, exhibiting significant differences from the non-private survival curves.

IV. DISCUSSION

We analyze the three time-to-event data privacy methods across eight different dimensions in Table IV. Combined with observations from the evaluation results, our discussion focuses on a few key aspects in Table IV, which we believe may help domain experts identify suitable privacy-protecting solutions.

We will first review the **input** and **output** of each privacy method. As can be seen, all privacy methods take the time-to-event dataset (i.e., D) as input. DPTime and BinSup output aggregate statistics over temporal intervals for each cohort, while TE-Sanitizer directly outputs a time-to-event dataset with individual records (i.e., $\hat{D}$). However, as done in our evaluation, a simple procedure can be performed to generate individual records according to the output aggregate statistics of DPTime or BinSup, where the time-to-event value in each generated record is rounded to the start of the temporal interval.

An important aspect for privacy-protecting healthcare data analytics is **data truthfulness**. In Table IV, we summarize both temporal truthfulness for each individual and cohort size truthfulness in the sanitized data. For temporal truthfulness, all the privacy methods studied in this manuscript modify the time-to-event values either by introducing random noise (i.e., TE-Sanitizer, DPTime) or by rounding to disjoint temporal intervals (i.e., BinSup). Among them, TE-Sanitizer and BinSup allow the data curator to have fine-grained control over the distortion for each time-to-event by tuning parameters (i.e., W and $TimeBin$). In our evaluation, we observe that DPTime and BinSup tend to inflict higher KL divergence (as well as MAE errors) in the sanitized time-to-events than TE-Sanitizer. For cohort size truthfulness, which represents how well the sanitized data preserves the number of individuals in each cohort, TE-Sanitizer does not change the number of individuals in each cohort. However, DPTime and BinSup may modify the number of individuals in each cohort, due to perturbed and suppressed aggregates, leading to cohort sizes different from real data. Furthermore, even with a low count threshold, the data loss inflicted by BinSup may significantly reduce the usability of the shared data, as shown in our Kaplan-Meier case studies.

Another aspect is the type of **privacy protection** provided these methods. TE-Sanitizer protects the time-to-event value under the definition of (ϵW)-TI. DPTime hides the presence of individual in the data under the ϵ-DP model. Both TE-Sanitizer and DPTime provide provable privacy guarantees. On the other hand, BinSup hides individuals via grouping and suppression, which may be more intuitive to practitioners.

The **empirical privacy risk** investigated in this manuscript shows that an adversary, who has knowledge about the target individual's time-to-event value, may successfully infer the cohort membership of the target. This risk level can be quite high in the non-private setting, as shown in Figure 1. Despite differences in the underlying privacy notion, all privacy methods show reduced empirical privacy risks in the sanitized dataset. The DPTime method, which satisfies differential privacy, has shown to achieve the lowest empirical privacy risk among all methods. We also observe that the BinSup method may provide an unbalanced privacy protection between cohorts, if the parameters are not carefully selected.

The input **data size** has effects on the empirical privacy risk in the non-private setting: our results show that the

privacy risk tends to decrease as the size of the dataset increases. The privacy methods considered in this work are effective in reducing the privacy risk, even with smaller data sizes. As the data size increases, the utility of the sanitize data increases for all privacy methods. TE-Sanitizer provides better utility measures at varying dataset sizes, greatly outperforming DP-Time and BinSup for smaller datasets.

While sharing **additional covariates** privately is out of the scope of this work, we provide our perspective briefly in the following. Very often, additional covariates are recorded along with the cohort label and time-to-event value for each individual, e.g., demographic information and pathology or test results. We have shown that the knowledge of the time-to-event value will increase the adversary's precision in inferring the individual's cohort label. It is a reasonable belief that releasing additional covariates about the individual will lead to a larger number of possible inference attacks (e.g., with knowledge about a combination of covariates) and thus higher empirical privacy risks. Recent privacy protecting solutions for sharing tabular data include applying randomized mechanisms to all attributes to achieve local differential privacy [38], [39] and synthesizing data records with generative adversarial networks [40], [41]. The former approach provides strong privacy guarantees for each attribute while introducing high perturbation cost. The latter leverages recent machine learning techniques to generate realistic synthetic data records; however, it is still an open question whether the synthetic data may be private and useful in specific analytics and applications [42].

V. Conclusion

In this work, we studied the applicability of privacy-protecting solutions for sharing health data with time-to-event values. We assessed the privacy-utility trade-off of traditional privacy solutions that protect data at aggregate-level (i.e., binning and differential privacy), and proposed a new method for individual-level time-to-event privacy protection. Overall, our study provides useful insights on the practical privacy risk of sharing time-to-event data, and aim to help domain experts (e.g., clinicians) identify privacy-protecting solutions that best suit their specific application settings.

Acknowledgment

The authors would like to thank the reviewers for valuable feedback. This work has been supported in part by National Human Genome Research Institute grant R00HG010493, NSF CNS-1951430, and UNC Charlotte. The opinions, findings, and conclusions or recommendations expressed in this material are those of the authors and do not necessarily reflect the views of the sponsors.

References

[1] X. Zuo, Y. Chen, L. Ohno-Machado, and H. Xu, "How do we share data in covid-19 research? a systematic review of covid-19 datasets in pubmed central articles," *Briefings in Bioinformatics*, vol. 22, no. 2, pp. 800–811, 2021.

[2] H. A. Rothan and S. N. Byrareddy, "The epidemiology and pathogenesis of coronavirus disease (covid-19) outbreak," *Journal of autoimmunity*, vol. 109, p. 102433, 2020.

[3] Y. Dong, X. Mo, Y. Hu, X. Qi, F. Jiang, Z. Jiang, and S. Tong, "Epidemiology of covid-19 among children in china," *Pediatrics*, vol. 145, no. 6, 2020.

[4] S. P. Adhikari, S. Meng, Y.-J. Wu, Y.-P. Mao, R.-X. Ye, Q.-Z. Wang, C. Sun, S. Sylvia, S. Rozelle, H. Raat *et al.*, "Epidemiology, causes, clinical manifestation and diagnosis, prevention and control of coronavirus disease (covid-19) during the early outbreak period: a scoping review," *Infectious diseases of poverty*, vol. 9, no. 1, pp. 1–12, 2020.

[5] B. Lau, S. R. Cole, and S. J. Gange, "Competing risk regression models for epidemiologic data," *American journal of epidemiology*, vol. 170, no. 2, pp. 244–256, 2009.

[6] P. C. Austin, D. S. Lee, and J. P. Fine, "Introduction to the analysis of survival data in the presence of competing risks," *Circulation*, vol. 133, no. 6, pp. 601–609, 2016.

[7] J. F. Tierney, L. A. Stewart, D. Ghersi, S. Burdett, and M. R. Sydes, "Practical methods for incorporating summary time-to-event data into meta-analysis," *Trials*, vol. 8, no. 1, pp. 1–16, 2007.

[8] M. Nemati, J. Ansary, and N. Nemati, "Machine-learning approaches in covid-19 survival analysis and discharge-time likelihood prediction using clinical data," *Patterns*, vol. 1, no. 5, p. 100074, 2020.

[9] P. Putzel, H. Do, A. Boyd, H. Zhong, and P. Smyth, "Dynamic survival analysis for ehr data with personalized parametric distributions," in *Proceedings of the 6th Machine Learning for Healthcare Conference*, ser. Proceedings of Machine Learning Research, K. Jung, S. Yeung, M. Sendak, M. Sjoding, and R. Ranganath, Eds., vol. 149. PMLR, 06–07 Aug 2021, pp. 648–673.

[10] L. Bonomi, X. Jiang, and L. Ohno-Machado, "Protecting patient privacy in survival analyses," *Journal of the American Medical Informatics Association*, vol. 27, no. 3, pp. 366–375, 2020.

[11] Y. Liu, C. Yan, Z. Yin, Z. Wan, W. Xia, M. Kantarcioglu, Y. Vorobeychik, E. W. Clayton, and B. A. Malin, "Biomedical research cohort membership disclosure on social media," in *AMIA Annual Symposium Proceedings*, vol. 2019. American Medical Informatics Association, 2019, p. 607.

[12] P. Umar, C. Akiti, A. Squicciarini, and S. Rajtmajer, "Self-disclosure on twitter during the covid-19 pandemic: A network perspective," in *Joint European Conference on Machine Learning and Knowledge Discovery in Databases*. Springer, 2021, pp. 271–286.

[13] Z. Lin, M. Hewett, and R. B. Altman, "Using binning to maintain confidentiality of medical data." in *Proceedings of the AMIA Symposium*. American Medical Informatics Association, 2002, p. 454.

[14] C. M. O'Keefe, R. S. Sparks, D. McAullay, and B. Loong, "Confidentialising survival analysis output in a remote data access system," *Journal of Privacy and Confidentiality*, vol. 4, no. 1, 2012.

[15] S. N. Murphy and H. C. Chueh, "A security architecture for query tools used to access large biomedical databases." in *Proceedings of the AMIA Symposium*. American Medical Informatics Association, 2002, p. 552.

[16] C. Dwork, A. Roth *et al.*, "The algorithmic foundations of differential privacy." *Foundations and Trends in Theoretical Computer Science*, vol. 9, no. 3-4, pp. 211–407, 2014.

[17] C. Curtis, S. P. Shah, S.-F. Chin, G. Turashvili, O. M. Rueda, M. J. Dunning, D. Speed, A. G. Lynch, S. Samarajiwa, Y. Yuan *et al.*, "The genomic and transcriptomic architecture of 2,000 breast tumours reveals novel subgroups," *Nature*, vol. 486, no. 7403, pp. 346–352, 2012.

[18] J. Domingo-Ferrer, "A survey of inference control methods for privacy-preserving data mining," in *Privacy-preserving data mining*. Springer, 2008, pp. 53–80.

[19] L. Willenborg and T. De Waal, *Elements of statistical disclosure control*. Springer Science & Business Media, 2012, vol. 155.

[20] R. Sparks, C. Carter, J. B. Donnelly, C. M. O'Keefe, J. Duncan, T. Keighley, and D. McAullay, "Remote access methods for exploratory data analysis and statistical modelling: Privacy-preserving analytics®," *Computer methods and programs in biomedicine*, vol. 91, no. 3, pp. 208–222, 2008.

[21] S. L. Warner, "Randomized response: A survey technique for eliminating evasive answer bias," *Journal of the American Statistical Association*, vol. 60, no. 309, pp. 63–69, 1965.

[22] J. C. Duchi, M. I. Jordan, and M. J. Wainwright, "Local privacy and statistical minimax rates," in *2013 IEEE 54th Annual Symposium on Foundations of Computer Science*. IEEE, 2013, pp. 429–438.

[23] B. Ding, J. Kulkarni, and S. Yekhanin, "Collecting telemetry data privately," *Advances in Neural Information Processing Systems*, vol. 30, 2017.

[24] Ú. Erlingsson, V. Pihur, and A. Korolova, "Rappor: Randomized aggregatable privacy-preserving ordinal response," in *Proceedings of the 2014 ACM SIGSAC conference on computer and communications security*, 2014, pp. 1054–1067.

[25] Q. Geng, P. Kairouz, S. Oh, and P. Viswanath, "The staircase mechanism in differential privacy," *IEEE Journal of Selected Topics in Signal Processing*, vol. 9, no. 7, pp. 1176–1184, 2015.

[26] K. Chatzikokolakis, M. E. Andrés, N. E. Bordenabe, and C. Palamidessi, "Broadening the scope of differential privacy using metrics," in *International Symposium on Privacy Enhancing Technologies Symposium*. Springer, 2013, pp. 82–102.

[27] M. E. Andrés, N. E. Bordenabe, K. Chatzikokolakis, and C. Palamidessi, "Geo-indistinguishability: Differential privacy for location-based systems," in *Proceedings of the 2013 ACM SIGSAC conference on Computer & communications security*, 2013, pp. 901–914.

[28] A. Ghosh, T. Roughgarden, and M. Sundararajan, "Universally utility-maximizing privacy mechanisms," *SIAM Journal on Computing*, vol. 41, no. 6, pp. 1673–1693, 2012.

[29] C. Dwork, F. McSherry, K. Nissim, and A. Smith, "Calibrating noise to sensitivity in private data analysis," in *Theory of cryptography conference*. Springer, 2006, pp. 265–284.

[30] D. Kifer and B.-R. Lin, "An axiomatic view of statistical privacy and utility," *Journal of Privacy and Confidentiality*, vol. 4, no. 1, 2012.

[31] L. Bonomi and L. Fan, "Sharing time-to-event data with privacy protection (with supplementary material)," https://bonomi-research.github.io/files/TE-full.pdf.

[32] B. Pereira, S.-F. Chin, O. M. Rueda, H.-K. M. Vollan, E. Provenzano, H. A. Bardwell, M. Pugh, L. Jones, R. Russell, S.-J. Sammut *et al.*, "The somatic mutation profiles of 2,433 breast cancers refine their genomic and transcriptomic landscapes," *Nature communications*, vol. 7, no. 1, pp. 1–16, 2016.

[33] B. Xu, B. Gutierrez, S. Mekaru, K. Sewalk, L. Goodwin, A. Loskill, E. L. Cohn, Y. Hswen, S. C. Hill, M. M. Cobo *et al.*, "Epidemiological data from the covid-19 outbreak, real-time case information," *Scientific data*, vol. 7, no. 1, pp. 1–6, 2020.

[34] N. Homer, S. Szelinger, M. Redman, D. Duggan, W. Tembe, J. Muehling, J. V. Pearson, D. A. Stephan, S. F. Nelson, and D. W. Craig, "Resolving individuals contributing trace amounts of dna to highly complex mixtures using high-density snp genotyping microarrays," *PLoS Genet*, vol. 4, no. 8, p. e1000167, 2008.

[35] N. Foldvary, B. Nashold, E. Mascha, E. Thompson, N. Lee, J. McNamara, D. Lewis, J. Luther, A. Friedman, and R. Radtke, "Seizure outcome after temporal lobectomy for temporal lobe epilepsy: a kaplan-meier survival analysis," *Neurology*, vol. 54, no. 3, pp. 630–630, 2000.

[36] J. M. Bland and D. G. Altman, "Survival probabilities (the kaplan-meier method)," *Bmj*, vol. 317, no. 7172, pp. 1572–1580, 1998.

[37] ——, "The logrank test," *Bmj*, vol. 328, no. 7447, p. 1073, 2004.

[38] X. Ren, C.-M. Yu, W. Yu, S. Yang, X. Yang, J. A. McCann, and P. S. Yu, "Lopub: High-dimensional crowdsourced data publication with local differential privacy," *IEEE Transactions on Information Forensics and Security*, vol. 13, no. 9, pp. 2151–2166, 2018.

[39] N. Wang, X. Xiao, Y. Yang, J. Zhao, S. C. Hui, H. Shin, J. Shin, and G. Yu, "Collecting and analyzing multidimensional data with local differential privacy," in *2019 IEEE 35th International Conference on Data Engineering (ICDE)*. IEEE, 2019, pp. 638–649.

[40] E. Choi, S. Biswal, B. Malin, J. Duke, W. F. Stewart, and J. Sun, "Generating multi-label discrete patient records using generative adversarial networks," in *Machine learning for healthcare conference*. PMLR, 2017, pp. 286–305.

[41] B. K. Beaulieu-Jones, Z. S. Wu, C. Williams, R. Lee, S. P. Bhavnani, J. B. Byrd, and C. S. Greene, "Privacy-preserving generative deep neural networks support clinical data sharing," *Circulation: Cardiovascular Quality and Outcomes*, vol. 12, no. 7, p. e005122, 2019.

[42] N. Ruiz, K. Muralidhar, and J. Domingo-Ferrer, "On the privacy guarantees of synthetic data: a reassessment from the maximum-knowledge attacker perspective," in *International Conference on Privacy in Statistical Databases*. Springer, 2018, pp. 59–74.