

Heterogeneous Datasets for Federated Survival Analysis Simulation

Alberto Archetti
alberto.archetti@polito.it
Politecnico di Milano
Milan, Italy

Eugenio Lomurno
eugenio.lomurno@polimi.it
Politecnico di Milano
Milan, Italy

Francesco Lattari
francesco.lattari@polimi.it
Politecnico di Milano
Milan, Italy

André Martin
andre.martin@tu-dresden.de
Technische Universität Dresden
Dresden, Germany

Matteo Matteucci
matteo.matteucci@polimi.it
Politecnico di Milano
Milan, Italy

ABSTRACT

Survival analysis studies time-modeling techniques for an event of interest occurring for a population. Survival analysis found widespread applications in healthcare, engineering, and social sciences. However, the data needed to train survival models are often distributed, incomplete, censored, and confidential. In this context, federated learning can be exploited to tremendously improve the quality of the models trained on distributed data while preserving user privacy. However, federated survival analysis is still in its early development, and there is no common benchmarking dataset to test federated survival models. This work provides a novel technique for constructing realistic heterogeneous datasets by starting from existing non-federated datasets in a reproducible way. Specifically, we propose two dataset-splitting algorithms based on the Dirichlet distribution to assign each data sample to a carefully chosen client: quantity-skewed splitting and label-skewed splitting. Furthermore, these algorithms allow for obtaining different levels of heterogeneity by changing a single hyperparameter. Finally, numerical experiments provide a quantitative evaluation of the heterogeneity level using log-rank tests and a qualitative analysis of the generated splits. The implementation of the proposed methods is publicly available in favor of reproducibility and to encourage common practices to simulate federated environments for survival analysis.

CCS CONCEPTS

• **Computing methodologies** → **Distributed computing methodologies**; • **Applied computing**;

KEYWORDS

datasets, federated learning, survival analysis

ACM Reference Format:

Alberto Archetti, Eugenio Lomurno, Francesco Lattari, André Martin, and Matteo Matteucci. 2023. Heterogeneous Datasets for Federated Survival Analysis Simulation. In *Companion of the 2023 ACM/SPEC International Conference on Performance Engineering (ICPE '23 Companion)*, April 15–19, 2023, Coimbra, Portugal. ACM, New York, NY, USA, 8 pages. <https://doi.org/10.1145/3578245.3584935>

1 INTRODUCTION

Survival analysis [21, 42] is a subfield of statistics focused on modeling the occurrence time of an event of interest for a population. In particular, its goal is to exploit statistical and machine learning techniques to provide a survival function, i.e., a function that estimates the event occurrence probability with respect to time for an individual. Survival analysis has been successfully applied in many healthcare, engineering, and social science applications [14]. However, the data to train survival models are often distributed, incomplete, inaccurate, and confidential [3, 36]. On top of that, survival data may include a considerable portion of censored observations, i.e., instances for which the event of interest has yet to occur. In censored samples, the observed time is an underestimation of the actual occurrence time of the event. As a result, data scarcity, censorship, and confidentiality can hinder the applicability of survival analysis when addressing real-world, large-scale problems.

In this context, Federated Learning (FL) [17, 26] holds tremendous potential to improve the effectiveness of survival analysis applications. FL is a subfield of distributed machine learning that investigates techniques to train machine learning models while preserving user privacy. In FL, data information never leaves the device in which it is produced, collected, and stored. FL allows for training on large-scale data, improving the quality, fairness, and generalizability of the resulting models with respect to the non-distributed counterparts.

Federated survival analysis studies the relationship between federated learning and survival analysis. In particular, survival models present structural components that make their inclusion into existing federated learning algorithms non-trivial [3, 29, 35, 43]. Since this field is in its early development, reproducible and standardized simulation environments are paramount for the comparability of results. Simulation environments mimic one or many aspects of real-world federations, such as client availability, communication constraints, computation constraints, and data heterogeneity. Some

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than the author(s) must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.
ICPE '23 Companion, April 15–19, 2023, Coimbra, Portugal

© 2023 Copyright held by the owner/author(s). Publication rights licensed to ACM.
ACM ISBN 979-8-4007-0072-9/23/04...\$15.00
<https://doi.org/10.1145/3578245.3584935>

existing works provide simulation environments for standard federated learning applications [5, 6]. However, these environments have no direct support for survival analysis problems. Other works implement algorithms for non-federated survival models [11, 15, 23, 33] based on centralized survival datasets [13]. Recently, a benchmarking suite for federated healthcare problems has been developed, including a single federated survival dataset with a predefined data split among 6 clients [39].

The goal of this work is to extend the benchmarking ground for federated survival models. To this end, we present a novel technique for constructing realistic federated datasets from existing non-federated survival datasets in a flexible and reproducible way. Realistic federated datasets mimic real-world heterogeneity by exhibiting non-identically distributed (non-IID) data among clients. More specifically, we provide two algorithms for assigning each data sample from a centralized survival dataset to a carefully chosen client. The proposed algorithms are based on the Dirichlet distribution [16, 25], as it can induce distribution skewness by tuning a single parameter. The first algorithm focuses on building federated datasets with a non-uniform number of samples. We call this algorithm quantity-skewed splitting. The second one, instead, builds client datasets with different label distributions. We call this algorithm label-skewed splitting. The heterogeneity level introduced by each algorithm in the resulting data assignments can be tuned with a parameter $\alpha > 0$, such that for $\alpha \rightarrow 0$ data are more skewed, while for $\alpha \rightarrow \infty$ data are more uniform. The ability to tune the heterogeneity level allows for federated simulations with different environmental conditions. This aspect is essential to test the resilience of federated survival models to non-IID realistic data distributions.

The presented techniques have been tested on a collection of datasets for survival analysis, providing visual insights about the level of heterogeneity induced in each setting. Also, the level of heterogeneity is numerically investigated with log-rank tests [7] within client distributions. The experimental evaluation demonstrates that the proposed techniques are able to build heterogeneous federated datasets starting from non-federated survival data. Moreover, the numerical analysis shows how the α parameter can effectively control the heterogeneity level induced by each split.

The implementation of quantity-skewed and label-skewed splitting is publicly available [4] in favor of reproducibility and to encourage the usage of common practices in the simulation of federated survival environments.

2 BACKGROUND AND RELATED WORKS

This section summarizes the main aspects of survival analysis and federated learning and reviews the state-of-the-art on federated survival analysis.

2.1 Survival Analysis

Survival analysis, also known as time-to-event analysis, is a statistical machine learning field that models the occurrence time of an event of interest for a population [42]. The distinctive feature of survival models is the handling of censored data. With censored data, we refer to samples for which the event occurrence was not

observed during the study. A survival dataset D is a set of N triplets

$$(\mathbf{x}_i, \delta_i, t_i), i = 1, \dots, N \text{ s.t.}$$

- $\mathbf{x}_i \in \mathbb{R}^d$ is a d -dimensional feature vector, also called co-variate vector, that retains all the input information for a sample;
- δ_i is the event occurrence indicator. If $\delta_i = 1$, then the i -th sample experienced the event, otherwise the i -th sample is censored and $\delta_i = 0$;
- $t_i = \min \{t_i^e, t_i^c\}$ is the minimum between the actual event time t_i^e and the censoring time t_i^c .

This setting refers to right-censoring [24], where the censoring time is less than or equal to the actual event time. This is the case, for instance, of disease recurrence under a certain treatment [38] or patient death [12]. Indeed, right-censoring is the most common scenario in real-world survival applications [42]. Therefore, we limit the discussion to the right censoring setting for the rest of the paper.

The goal of survival analysis is to estimate the event occurrence probability with respect to time. In particular, the output of a survival model is the survival function

$$S(t|\mathbf{x}) = P(T > t|\mathbf{x}).$$

Survival models are classified into three types: non-parametric, semi-parametric, and parametric [42]. In this work, we include non-parametric models in the analysis of the proposed data splitting algorithms, as these are the only models that make no assumption about the underlying event distribution over time. Moreover, non-parametric models are well-suited for survival data visualization. Indeed, non-parametric models encode the overall survival behavior of a population by predicting a survival function $\hat{S}(t)$ which is not conditioned on \mathbf{x} .

Non-parametric models are Kaplan-Meier (KM) [18], Nelson-Aalen [1, 32], and Life-Table [10]. Among those, the KM estimator is the most widely spread in survival applications due to its intuitive interpretation. The KM estimator starts from the set of unique event occurrence times $T_D = \{t_j : (\mathbf{x}_i, \delta_i, t_j) \in D\}$. Then, for each $t_j \in T_D$ it computes the number of observed events $d_j \geq 1$ at time t_j and the number of samples r_j that did not yet experience an event. The KM estimator is computed as

$$\hat{S}(t) = \prod_{j:t_j < t} \left(1 - \frac{d_j}{r_j}\right).$$

2.2 Federated Learning

Federated Learning (FL) [17, 26] is a machine learning setting in which a set of agents jointly train a model without sharing the data they store locally. FL algorithms rely on a central server for message exchange and agent coordination. A federation is composed of K clients, each holding a private dataset D_k , $k = 1, \dots, K$. The goal of a FL algorithm is to find the best parameters \mathbf{w} that optimize a global loss function \mathcal{L} :

$$\min_{\mathbf{w}} \mathcal{L}(\mathbf{w}) = \min_{\mathbf{w}} \sum_{k=1}^K \lambda_k \mathcal{L}_k(\mathbf{w}).$$

\mathcal{L}_k is the local loss function computed by client k . λ_k is a set of parameters weighting the contribution of each client to the global

loss. Usually, λ_k is proportional to the number of samples on which each client k evaluated $\mathcal{L}_k(w)$ locally. This weighting strategy favors contributions from clients holding more private data, which are more likely to be representative of the entire data distribution.

Federated Averaging (FedAvg) [30] is the first algorithm developed to minimize \mathcal{L} . It relies on iterative averaging of model parameters trained locally on random subsets of clients. However, FedAvg is not always suited to face system security and confidentiality preservation challenges in real-world applications [31, 34]. Moreover, real-world applications present multiple levels of heterogeneity. First, system heterogeneity constraints FL algorithms to comply with the hardware limitations of the network channel and the clients' devices. Second, datasets are not guaranteed to contain identically distributed data. In fact, in most real-world scenarios data are likely to be non-IID. In order to handle data heterogeneity in federated environments, several non-survival federated algorithms have been proposed [2, 19, 27].

2.3 Federated Survival Analysis

Federated learning provides key advantages for the future of healthcare applications [36]. In particular, federated survival analysis investigates the opportunities and challenges related to the integration of federated learning into survival analysis tasks. However, few works specifically tackle federated survival analysis applications. Some works [3, 43] provide solutions for the non-separability of the partial log-likelihood loss, used to train Cox survival models [9]. Indeed, non-separable loss functions are not suited for federated learning algorithms, as their evaluation requires access to all the available data in the federation. Other works [29, 35] provide federated versions of classical survival algorithms asymptotically equivalent to their centralized counterparts. Within these works, data federations are built with uniform data splits or with entirely simulated datasets.

2.4 Federated Datasets

Concerning the available datasets for federated simulation, LEAF [8] is the most widely spread dataset collection for standard federated learning applications. It provides several real-world datasets covering classification, sentiment analysis, next-character, and next-word prediction. Secure Generative Data Exchange (SGDE) [28] is a recent framework to build synthetic datasets in a privacy-preserving way. SGDE provides inherently heterogeneous datasets composed of synthetic samples provided by client-side data generators. Currently, SGDE has been applied to classification and regression problems only. Other studies [16, 25] investigate the taxonomy of data heterogeneity and provide techniques to emulate non-IID data splits starting from centralized classification datasets. Recently, FLamby [39] provided a set of benchmarking tools for cross-silo federated applications concerning healthcare. Among those, Fed-TCGA-BRCA is a federated survival dataset collecting the data of 1066 patients geographically grouped into 6 clients.

To the best of our knowledge, Fed-TCGA-BRCA is the only federated survival dataset proposed to date. Moreover, existing data-splitting techniques are tailored for non-survival problems only.

This is the first study extending data-splitting techniques to survival analysis, providing greater flexibility for modeling simulated survival environments.

3 METHOD

This paper presents a set of techniques to split survival datasets into heterogeneous federations. We start from a survival dataset D and a number of clients K . The goal is to assign to each sample in D a client $k \in \{1, \dots, K\}$, such that federated survival algorithms can leverage the set of D_k s to simulate heterogeneous learning scenarios. The work proposes two splitting techniques: quantity-skewed and label-skewed splitting.

3.1 Quantity-Skewed Splitting

Quantity-skewed splitting pertains to a scenario where the number of samples for each client k , represented as $|D_k|$, varies among clients. In such a scenario, clients with a limited number of samples may generate gradients that are inherently noisy, which can impede the convergence of federated learning algorithms. This is due to the fact that clients with a smaller number of samples tend to exhibit higher variance in their gradients, leading to instability in the federated learning process and hampering convergence rate.

Simulation of quantity-skewed scenarios is essential in assessing the robustness of federated survival algorithms. It enables researchers to evaluate the algorithm's ability to handle the imbalance in sample distribution across clients and its impact on algorithm performance.

Similarly to [16, 25], the proportion of samples \mathbf{p} to assign to each client follows a Dirichlet distribution

$$\mathbf{p} \sim \mathcal{D}(\alpha \cdot \mathbf{1}_K).$$

Here, $\mathbf{1}_K$ is a vector of 1s of length K . $\mathbf{p} \in [0, 1]^K$ such that $\langle \mathbf{1}_K, \mathbf{p} \rangle = 1$. $\alpha > 0$ is a similarity parameter controlling the similarity between client dataset cardinalities $|D_k|$. For $\alpha \rightarrow 0$, the number of samples for each D_k are heterogeneous. Conversely, for $\alpha \rightarrow \infty$, the number of samples for each D_k tends to be similar. With quantity-skewed splitting, each sample $(\mathbf{x}_i, \delta_i, t_i)$ is assigned to a client dataset D_k with probability

$$P((\mathbf{x}_i, \delta_i, t_i) \in D_k) = \mathbf{p}[k].$$

3.2 Label-Skewed Splitting

Label-skewed splitting pertains to scenarios in which the distribution of labels differs among client datasets. This type of distribution heterogeneity is commonly encountered in real-world federated learning scenarios. The non-IID distribution can be attributed to various factors, including variations in data collection and storage processes, the use of different acquisition devices, and variations in preprocessing or labeling techniques. Additionally, clients may have different label quantities due to domain-specific factors. For instance, in a federated healthcare scenario for treatment risk assessment, one client may have a dataset of records from a rural hospital, while another client may have data from an urban hospital. These datasets from different locations may exhibit heterogeneous label distributions due to disparities in patient demographics and healthcare access.

Table 1: Survival datasets involved in the experiments.

Dataset	Samples	Censored	Features
GBSG [38]	686	44%	8
METABRIC [20]	1904	58%	8
AIDS [37]	2839	62%	4
FLCHAIN [40]	7874	28%	10
SUPPORT [41]	9105	68%	35

To produce a label-skewed data split, first, the timeline of the original survival dataset is divided into B bins, obtaining a set of time instants $\{\tau_0, \dots, \tau_B\}$. The bin identification can be uniform or quantile-based, as in [22]. Then, each sample $(\mathbf{x}_i, \delta_i, t_i)$ is assigned a class that corresponds to the b -th bin, such that $t_i \in (\tau_{b-1}, \tau_b]$. Following [16, 25], the Dirichlet distribution is used to identify heterogeneous splitting proportions according to the sample class as

$$\begin{aligned} \mathbf{p}_1 &\sim \mathcal{D}(\alpha \cdot \mathbf{1}_K) \\ &\vdots \\ \mathbf{p}_B &\sim \mathcal{D}(\alpha \cdot \mathbf{1}_K) \end{aligned}$$

Finally, each sample $(\mathbf{x}_i, \delta_i, t_i)$ assigned to label b is added to D_k with probability

$$P((\mathbf{x}_i, \delta_i, t_i) \in D_k) = \mathbf{p}_b[k].$$

The α parameter controls the level of similarity between label distributions. For $\alpha \rightarrow \infty$, client label distributions are similar, while for $\alpha \rightarrow 0$ label distributions differ. The numerical dependency between α and the data heterogeneity level is discussed in detail using log-rank tests [7] in Section 4.

4 EXPERIMENTS

This section presents the experiments carried out to evaluate the proposed methods for building heterogeneous datasets for federated survival analysis.

4.1 Datasets

Each of the experiments involves the following survival datasets: the German Breast Cancer Study Group 2 (GBSG2) [38], the Molecular Taxonomy of Breast Cancer International Consortium (METABRIC) [20], the Australian AIDS survival dataset (AIDS) [37], the assay of serum-free light chain dataset (FLCHAIN) [40], and the Study to Understand Prognoses Preferences Outcomes and Risks of Treatment (SUPPORT) [41]. The dataset summary statistics are collected in Table 1.

4.2 Visualizing Splitting Methods

This section describes the visual results obtained from the splitting methods under different α parameters. In particular, Figure 1 shows the results of the quantity-skewed splitting algorithm described in Section 3.1. Splits are generated for a federation of 10 clients ($K = 10$), each identified by a different color in the plots. In Figure 1, each row corresponds to one of the example datasets described in Section 4.1. Columns, instead, refer to different values of the similarity parameter α , ranging from $\alpha = 1000$ (low heterogeneity)

to $\alpha = 0.5$ (high heterogeneity). Each plot shows the client dataset cardinalities $|D_k|$ with respect to clients $k = 1, \dots, 10$. By looking at the plots on the left of Figure 1, higher values of α tend to produce similar dataset cardinalities $|D_k|$. Conversely, for lower α values, $|D_k|$ s considerably differ between clients. This trend is confirmed for all the datasets involved in the experiments.

Figure 2 shows the results of the label-skewed splitting algorithm described in Section 3.2. Each plot shows the Kaplan-Meier estimators $\hat{S}_k(t)$ of each client dataset D_k , $k = 1, \dots, 10$. The KM estimator an excellent tool for survival function visualization, as it encodes the summary information concerning the survival labels in the dataset. Following the structure of Figure 1, in Figure 2 each row corresponds to a specific dataset from Section 4.1 and each column corresponds to an α value, decreasing from 1000 to 0.5. From the left column to the right column, the KM estimators of each client tend to separate, as α decreases. This is expected, as for lower α values, the Dirichlet distribution assigns non-uniform proportions of samples from each time bin to the clients. In this way, decreasing the α parameter results in heterogeneous federated distributions.

4.3 Numerical Analysis of Heterogeneity

This section provides the quantitative analysis carried out to evaluate the level of heterogeneity induced by each splitting method. A high level of data heterogeneity entails different client data distributions, which leads to more realistic federations. We use the log-rank test [7] to determine whether the event occurrence distribution is the same for two clients. This test verifies the null hypothesis that there is no statistically significant difference between the survival distributions of two given populations. We consider the distribution difference between two clients k_1, k_2 statistically significant if the p-value p_{k_1, k_2} resulting from the test is ≤ 0.05 .

In order to summarize the results for a federation, we define the heterogeneity score h of a federation as the fraction of client pairs $\mathcal{P} = \{(k_1, k_2) : k_1 < k_2 \wedge k_1, k_2 = 1, \dots, K\}$ for which the distribution difference is statistically significant, i.e.,

$$h = \frac{1}{|\mathcal{P}|} \sum_{(k_1, k_2) \in \mathcal{P}} \mathbb{1}(p_{k_1, k_2} \leq 0.05).$$

Table 2 collects the h values for quantity-skewed and label-skewed splits under several K and α values. Each result is averaged over 100 runs.

Concerning quantity-skewed splitting, each setting presents an average heterogeneity score smaller than 5%. In other words, quantity-skewed survival data does not present statistically significant label distribution differences when comparing pairs of client datasets. This implies that quantity-skewed splitting is well suited to test the resilience of a federated algorithm to data imbalance, but not to heterogeneous data distributions.

Conversely, label-skewed splitting exhibits noticeable differences in h scores depending on the value of α . In fact, for all the tested datasets, the h score with $\alpha = 1000$ is almost zero, and decreasing α affects the number of different label distributions among clients. For datasets with smaller total cardinalities (GBSG2, METABRIC, and AIDS) α must be smaller than 10 in order to detect noticeable differences between client distributions. Instead, datasets with more total samples (FLCHAIN and SUPPORT) present high heterogeneity



Figure 1: Number of samples $|D_k|$ for each client $k = 1, \dots, 10$. Each row refers to one of the datasets described in Section 4.1. Each column corresponds to a quantity-skewed split (Section 3.1) with a fixed similarity parameter α .

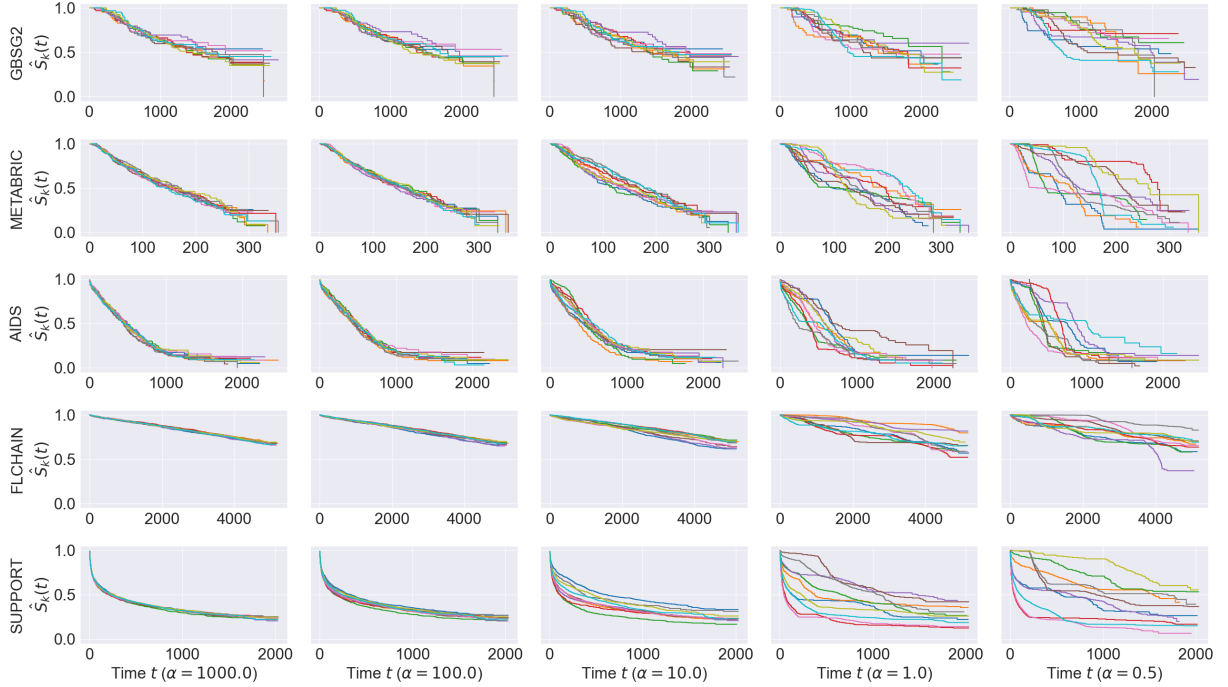


Figure 2: Kaplan-Meier estimators $\hat{S}_k(t)$ for each client $k = 1, \dots, 10$. Each row refers to one of the datasets described in Section 4.1. Each column corresponds to a label-skewed split (Section 3.2) with a fixed similarity parameter α .

Table 2: Heterogeneity score h for several clients K and similarity parameters α . Heterogeneity scores h are averaged over 100 runs and scaled by 100 for better readability.

Quantity-Skewed Split, $K = 5$						
Dataset	$\alpha = 1000.0$	$\alpha = 100.0$	$\alpha = 10.0$	$\alpha = 1.0$	$\alpha = 0.5$	$\alpha = 0.1$
GBSG2	2.6±6.0	3.1±7.1	3.4±8.2	2.1±5.9	4.3±9.3	2.2±6.6
METABRIC	2.8±7.3	3.3±7.9	3.1±7.6	2.9±8.3	1.5±5.4	2.2±7.5
AIDS	1.4±5.3	2.8±6.5	2.1±5.0	4.6±10.5	4.6±9.8	2.3±5.8
FLCHAIN	1.9±4.6	3.2±6.9	2.3±6.0	3.8±8.4	2.9±9.8	2.6±6.8
SUPPORT	3.0±6.9	2.0±4.7	2.5±6.7	3.3±7.4	3.7±9.4	0.3±2.2
Quantity-Skewed Split, $K = 10$						
Dataset	$\alpha = 1000.0$	$\alpha = 100.0$	$\alpha = 10.0$	$\alpha = 1.0$	$\alpha = 0.5$	$\alpha = 0.1$
GBSG2	4.1±4.8	3.9±5.0	3.0±4.9	3.0±4.3	3.0±4.5	1.9±3.3
METABRIC	3.6±5.3	4.7±6.0	4.4±5.9	3.5±5.6	3.6±4.7	1.7±4.0
AIDS	4.1±5.6	4.5±5.5	3.8±5.4	4.5±6.4	4.5±6.3	2.3±3.6
FLCHAIN	3.7±4.8	3.4±5.0	3.6±4.5	5.5±6.7	4.2±6.2	2.3±4.0
SUPPORT	4.1±5.8	3.4±4.6	4.0±4.8	3.9±5.7	4.2±6.5	1.0±2.3
Quantity-Skewed Split, $K = 50$						
Dataset	$\alpha = 1000.0$	$\alpha = 100.0$	$\alpha = 10.0$	$\alpha = 1.0$	$\alpha = 0.5$	$\alpha = 0.1$
GBSG2	3.9±2.0	3.4±1.8	3.5±2.0	3.0±1.8	2.6±1.7	1.6±1.0
METABRIC	4.6±2.3	4.7±2.6	4.4±2.0	3.9±2.1	3.2±1.8	1.5±1.1
AIDS	4.5±2.2	4.9±2.5	4.4±2.1	4.6±2.4	4.2±2.4	2.0±1.1
FLCHAIN	4.8±2.4	5.0±2.4	4.6±2.2	4.7±2.6	4.8±2.7	2.1±1.5
SUPPORT	4.5±2.2	4.5±2.3	4.8±2.3	3.9±2.1	3.4±2.1	0.8±0.8
Label-Skewed Split, $K = 5$						
Dataset	$\alpha = 1000.0$	$\alpha = 100.0$	$\alpha = 10.0$	$\alpha = 1.0$	$\alpha = 0.5$	$\alpha = 0.1$
GBSG2	0.2±2.0	0.1±1.0	5.8±9.4	46.7±20.9	58.2±17.0	73.8±18.2
METABRIC	0.0±0.0	0.5±2.2	20.9±17.2	66.1±19.0	76.7±14.5	82.3±13.3
AIDS	0.3±1.7	3.1±7.2	37.5±21.9	75.1±16.2	81.5±14.4	86.6±11.3
FLCHAIN	0.3±1.7	12.6±14.9	58.8±17.6	83.9±12.4	88.0±11.4	94.1±7.0
SUPPORT	0.5±2.2	29.6±20.8	74.3±15.7	91.3±9.7	92.5±7.4	94.0±6.4
Label-Skewed Split, $K = 10$						
Dataset	$\alpha = 1000.0$	$\alpha = 100.0$	$\alpha = 10.0$	$\alpha = 1.0$	$\alpha = 0.5$	$\alpha = 0.1$
GBSG2	0.4±1.5	0.6±1.5	2.8±4.3	32.2±11.7	43.7±11.5	63.2±12.9
METABRIC	0.1±0.4	0.2±1.0	10.6±8.4	54.6±13.6	66.5±10.1	76.7±8.7
AIDS	0.3±1.0	1.4±2.7	24.7±12.6	68.1±9.0	74.0±9.1	77.7±8.4
FLCHAIN	0.4±1.2	4.2±5.5	42.8±13.0	78.2±8.8	84.9±5.6	89.3±5.8
SUPPORT	0.1±0.4	14.7±9.7	63.2±10.5	87.0±4.9	88.5±4.7	89.7±6.1
Label-Skewed Split, $K = 50$						
Dataset	$\alpha = 1000.0$	$\alpha = 100.0$	$\alpha = 10.0$	$\alpha = 1.0$	$\alpha = 0.5$	$\alpha = 0.1$
GBSG2	0.5±0.6	0.6±0.6	0.5±0.6	5.7±2.2	10.8±3.3	23.8±4.2
METABRIC	0.2±0.3	0.3±0.5	1.3±1.2	21.7±4.5	33.1±5.2	48.8±5.5
AIDS	0.6±0.5	0.8±0.8	4.5±2.3	34.6±5.2	45.3±4.4	49.8±4.9
FLCHAIN	0.2±0.3	0.6±0.6	10.6±3.7	55.4±4.1	64.8±2.6	72.4±3.6
SUPPORT	0.0±0.0	0.5±0.6	29.0±5.5	69.9±2.7	75.5±2.3	73.4±4.2

even for $\alpha = 100$. For all the dataset splits in small federations ($K = 5$ and $K = 10$), α values smaller than 1 result in $h > 50\%$. The trend

does not apply to federations with more clients ($K = 50$), where even $\alpha = 0.1$ is not enough to obtain $h > 50\%$.

5 CONCLUSION

This paper proposes two algorithms to simulate data heterogeneity in survival datasets for federated learning. Federated simulation is an important step in survival analysis toward the implementation and production of more accurate, fair, and privacy-preserving survival models. The presented splitting techniques are based on the Dirichlet distribution. Quantity-skewed splitting produces datasets with variable cardinalities, while label-skewed splitting relies on time binning to split samples according to different label distributions. Visual insights are provided to show the behavior of the proposed methods under hyperparameter change. Moreover, log-rank tests are reported to provide a quantitative evaluation of the degree of heterogeneity induced by each data split. To encourage the adoption of common benchmarking practices for future experiments on federated survival analysis, we make the source code of the proposed algorithms publicly available.

ACKNOWLEDGMENTS

This project has been supported by AI-SPRINT: AI in Secure Privacy-preserving computing continuum (European Union's H2020 grant agreement No 101016577).

REFERENCES

- [1] Odd Aalen. 1978. Nonparametric inference for a family of counting processes. *The Annals of Statistics* (1978), 701–726.
- [2] Durmus Alp Emre Acar, Yue Zhao, Ramon Matas Navarro, Matthew Mattina, Paul N Whatmough, and Venkatesh Saligrama. 2021. Federated learning based on dynamic regularization. *arXiv preprint arXiv:2111.04263* (2021).
- [3] Mathieu Andreux, Andre Manoel, Romuald Menuet, Charlie Saillard, and Chloé Simpson. 2020. Federated survival analysis with discrete-time Cox models. *arXiv preprint arXiv:2006.08997* (2020).
- [4] Alberto Archetti, Eugenio Lomurno, Lattari Francesco, André Martin, and Matteo Matteucci. 2023. Heterogeneous Datasets for Federated Survival Analysis Simulation. <https://doi.org/10.5281/zenodo.7661027>
- [5] The TensorFlow Federated Authors. 2018. *TensorFlow Federated*. <https://github.com/tensorflow/federated>
- [6] Daniel J Beutel, Taner Topal, Akhil Mathur, Xinchu Qiu, Titouan Parcollet, Pedro PB de Gusmão, and Nicholas D Lane. 2020. Flower: A friendly federated learning research framework. *arXiv preprint arXiv:2007.14390* (2020).
- [7] J Martin Bland and Douglas G Altman. 2004. The logrank test. *Bmj* 328, 7447 (2004), 1073.
- [8] Sebastian Caldas, Sai Meher Karthik Duddu, Peter Wu, Tian Li, Jakub Konečný, H Brendan McMahan, Virginia Smith, and Ameet Talwalkar. 2018. Leaf: A benchmark for federated settings. *arXiv preprint arXiv:1812.01097* (2018).
- [9] D. R. Cox. 1972. Regression Models and Life-Tables. *Journal of the Royal Statistical Society. Series B (Methodological)* 34, 2 (1972), 187–220. <http://www.jstor.org/stable/2985181>
- [10] Sidney J Cutler and Fred Ederer. 1958. Maximum utilization of the life table method in analyzing survival. *Journal of chronic diseases* 8, 6 (1958), 699–712.
- [11] Cameron Davidson-Pilon. 2019. lifelines: survival analysis in Python. *Journal of Open Source Software* 4, 40 (2019), 1317.
- [12] Angela Dispenzieri, Jerry A Katzmann, Robert A Kyle, Dirk R Larson, Terry M Therneau, Colin L Colby, Raynell J Clark, Graham P Mead, Shaji Kumar, L Joseph Melton III, et al. 2012. Use of nonclonal serum immunoglobulin free light chains to predict overall survival in the general population. In *Mayo Clinic Proceedings*, Vol. 87. Elsevier, 517–523.
- [13] Erik Drysdale. 2022. SurvSet: An open-source time-to-event dataset repository. *arXiv preprint arXiv:2203.03094* (2022).
- [14] Frank Emmert-Streib and Matthias Dehmer. 2019. Introduction to survival analysis in practice. *Machine Learning and Knowledge Extraction* 1, 3 (2019), 1013–1038.
- [15] Stephane Fotso et al. 2019–. PySurvival: Open source package for Survival Analysis modeling. <https://www.pyssurvival.io/>
- [16] Tzu-Ming Harry Hsu, Hang Qi, and Matthew Brown. 2019. Measuring the effects of non-identical data distribution for federated visual classification. *arXiv preprint arXiv:1909.06335* (2019).
- [17] Peter Kairouz, H Brendan McMahan, Brendan Avent, Aurélien Bellet, Mehdi Benini, Arjun Nitin Bhagoji, Kallista Bonawitz, Zachary Charles, Graham Cormode, Rachel Cummings, et al. 2021. Advances and open problems in federated learning. *Foundations and Trends in Machine Learning* 14, 1–2 (2021), 1–210.
- [18] Edward L Kaplan and Paul Meier. 1958. Nonparametric estimation from incomplete observations. *Journal of the American statistical association* 53, 282 (1958), 457–481.
- [19] Sai Praneeth Karimireddy, Satyen Kale, Mehryar Mohri, Sashank Reddi, Sebastian Stich, and Ananda Theertha Suresh. 2020. Scaffold: Stochastic controlled averaging for federated learning. In *International Conference on Machine Learning*. PMLR, 5132–5143.
- [20] Jared L Katzman, Uri Shaham, Alexander Cloninger, Jonathan Bates, Tingting Jiang, and Yuval Kluger. 2018. DeepSurv: personalized treatment recommender system using a Cox proportional hazards deep neural network. *BMC medical research methodology* 18, 1 (2018), 1–12.
- [21] John P Klein and Melvin L Moeschberger. 2003. *Survival analysis: techniques for censored and truncated data*. Vol. 1230. Springer.
- [22] Håvard Kvamme and Ørnulf Borgan. 2021. Continuous and discrete-time survival prediction with neural networks. *Lifetime Data Analysis* 27, 4 (2021), 710–736.
- [23] Håvard Kvamme, Ørnulf Borgan, and Ida Scheel. 2019. Time-to-event prediction with neural networks and Cox regression. *arXiv preprint arXiv:1907.00825* (2019).
- [24] Elisa T Lee and John Wang. 2003. *Statistical methods for survival data analysis*. Vol. 476. John Wiley & Sons.
- [25] Qinqin Li, Yiqun Diao, Quan Chen, and Bingsheng He. 2022. Federated learning on non-iid data silos: An experimental study. In *2022 IEEE 38th International Conference on Data Engineering (ICDE)*. IEEE, 965–978.
- [26] Tian Li, Anit Kumar Sahu, Ameet Talwalkar, and Virginia Smith. 2020. Federated learning: Challenges, methods, and future directions. *IEEE Signal Processing Magazine* 37, 3 (2020), 50–60.
- [27] Tian Li, Anit Kumar Sahu, Manzil Zaheer, Maziar Sanjabi, Ameet Talwalkar, and Virginia Smith. 2020. Federated optimization in heterogeneous networks. *Proceedings of Machine Learning and Systems* 2 (2020), 429–450.
- [28] Eugenio Lomurno, Alberto Archetti, Lorenzo Cazzella, Stefano Samele, Leonardo Di Perna, and Matteo Matteucci. 2022. SGDE: Secure Generative Data Exchange for Cross-Silo Federated Learning. In *AIPR 2022, International Conference on Artificial Intelligence and Pattern Recognition*.
- [29] Chia-Lun Lu, Shuang Wang, Zhanglong Ji, Yuan Wu, Li Xiong, Xiaoqian Jiang, and Lucila Ohno-Machado. 2015. WebDISCO: a web service for distributed cox model learning without patient-level data sharing. *Journal of the American Medical Association* 312, 6 (2015), 1212–1219.
- [30] Brendan McMahan, Eider Moore, Daniel Ramage, Seth Hampson, and Blaise Agüera y Arcas. 2017. Communication-efficient learning of deep networks from decentralized data. In *Artificial intelligence and statistics*. PMLR, 1273–1282.
- [31] Viraaji Mothukuri, Reza M Parizi, Seyedamin Pouriyeh, Yan Huang, Ali Dehghantaha, and Gautam Srivastava. 2021. A survey on security and privacy of federated learning. *Future Generation Computer Systems* 115 (2021), 619–640.
- [32] Wayne Nelson. 1972. Theory and applications of hazard plotting for censored failure data. *Technometrics* 14, 4 (1972), 945–966.
- [33] Sebastian Pölsterl. 2020. scikit-survival: A Library for Time-to-Event Analysis Built on Top of scikit-learn. *Journal of Machine Learning Research* 21, 212 (2020), 1–6. <http://jmlr.org/papers/v21/p21-20-729.html>
- [34] Shadi Rahimian, Raouf Kerkouche, Ina Kurth, and Mario Fritz. 2022. Practical Challenges in Differentially-Private Federated Survival Analysis of Medical Data. In *Conference on Health, Inference, and Learning*. PMLR, 411–425.
- [35] Md Mahmudur Rahman and Sanjay Purushotham. 2022. FedPseudo: Pseudo value-based Deep Learning Models for Federated Survival Analysis. *arXiv preprint arXiv:2207.05247* (2022).
- [36] Nicola Rieke, Jonny Hancox, Wenqi Li, Fausto Milletari, Holger R Roth, Shadi Albarqouni, Spyridon Bakas, Mathieu N Galtier, Bennett A Landman, Klaus Maier-Hein, et al. 2020. The future of digital health with federated learning. *NPJ digital medicine* 3, 1 (2020), 1–7.
- [37] Brian Ripley, Bill Venables, Douglas M. Bates, Kurt Hornik, Albrecht Gebhardt, and David Firth. 2022. R package: MASS. <https://stat.ethz.ch/R-manual/R-devel/library/MASS/html/00Index.html>
- [38] M Schumacher, G Bastert, H Bojar, K Hübner, M Olschewski, W Sauerbrei, C Schmoor, C Beyerle, RL Neumann, and HF Rauschecker. 1994. Randomized 2 x 2 trial evaluating hormonal treatment and the duration of chemotherapy in node-positive breast cancer patients. German Breast Cancer Study Group. *Journal of Clinical Oncology* 12, 10 (1994), 2086–2093.
- [39] Jean Ogier du Terrail, Samy-Safwan Ayed, Edwige Cyffers, Felix Grimberg, Chaoyang He, Regis Loeb, Paul Mangold, Tanguy Marchand, Othmane Marfoq, Erum Mushtaq, et al. 2022. FLamby: Datasets and Benchmarks for Cross-Silo Federated Learning in Realistic Healthcare Settings. *arXiv preprint arXiv:2210.04620* (2022).
- [40] Terry Therneau, Thomas Lumley, Elizabeth Atkinson, and Cynthia Crowson. 2023. R package: survival. <https://stat.ethz.ch/R-manual/R-devel/library/survival/html/00Index.html>
- [41] Vanderbilt University Department of Biostatistics. 2022. Vanderbilt Biostatistics Datasets. <http://hbiostat.org/data>
- [42] Ping Wang, Yan Li, and Chandan K Reddy. 2019. Machine learning for survival analysis: A survey. *ACM Computing Surveys (CSUR)* 51, 6 (2019), 1–36.

[43] Xuan Wang, Harrison G Zhang, Xin Xiong, Chuan Hong, Griffin M Weber, Gabriel A Brat, Clara-Lea Bonzel, Yuan Luo, Rui Duan, Nathan P Palmer, et al.

2022. SurvMaximin: robust federated approach to transporting survival risk prediction models. *Journal of biomedical informatics* 134 (2022), 104176.