

Distributed Statistical Analyses: A Scoping Review and Examples of Operational Frameworks Adapted to Health Analytics

Félix Camirand Lemyre, Simon Lévesque, Marie-Pier Domingue, Klaus Herrmann, Jean-François Ethier

Submitted to: JMIR Medical Informatics
on: October 12, 2023

Disclaimer: © The authors. All rights reserved. This is a privileged document currently under peer-review/community review. Authors have provided JMIR Publications with an exclusive license to publish this preprint on its website for review purposes only. While the final peer-reviewed paper may be licensed under a CC BY license on publication, at this stage authors and publisher expressly prohibit redistribution of this draft paper other than for review purposes.

Table of Contents

Original Manuscript.....	5
Supplementary Files.....	50
Figures	51
Figure 1.....	52
Figure 2.....	53
Figure 3.....	54
Figure 4.....	55
Figure 5.....	56
Multimedia Appendixes	57
Multimedia Appendix 1.....	58
Multimedia Appendix 2.....	58
CONSORT (or other) checklists.....	59
CONSORT (or other) checklist 3.....	59

Distributed Statistical Analyses: A Scoping Review and Examples of Operational Frameworks Adapted to Health Analytics

Félix Camirand Lemyre^{1, 2*} PhD; Simon Lévesque^{1, 3, 2*} BSc; Marie-Pier Domingue^{1, 4} BSc; Klaus Herrmann² PhD; Jean-François Ethier^{1, 3, 5} MD, PhD

¹GRIIS Université de Sherbrooke Sherbrooke CA

²Département de mathématiques Faculté des sciences Université de Sherbrooke Sherbrooke CA

³Health Data Research Network Canada Vancouver CA

⁴Chaire MEIE Québec - Le numérique au service des systèmes de santé apprenants Université de Sherbrooke Sherbrooke CA

⁵Département de médecine Faculté de médecine et des sciences de la santé Université de Sherbrooke Sherbrooke CA

*these authors contributed equally

Corresponding Author:

Jean-François Ethier MD, PhD

GRIIS

Université de Sherbrooke

2500, boul. de l'Université

Sherbrooke

CA

Abstract

Background: Data from multiple organizations are crucial for advancing learning health systems. However, ethical, legal, and social concerns may restrict the use of standard statistical methods that rely on pooling data. Although distributed algorithms offer alternatives, they may not always be suitable for healthcare research frameworks.

Objective: This paper aims to support researchers and data custodians in three ways: (1) providing a concise overview of the literature on statistical inference methods for horizontally partitioned data; (2) describing the methods applicable to generalized linear models (GLM) and assessing their underlying distributional assumptions; (3) adapting existing methods to make them fully usable in healthcare research.

Methods: A scoping review methodology was employed for the literature mapping using interdisciplinary databases: MEDLINE, Scopus, MathSciNet and zbMATH. Included articles had to provide a methodological contribution and address inferential statistics on horizontally partitioned data. Model type, methodological setting and number of communications required between data nodes and the coordinating centre were among the data extracted. The methods presenting a methodological framework for GLM analyses with horizontally partitioned data were identified and assessed from the perspective of applicability in healthcare research. Statistical theory was used to adapt methods and derive the properties of the resulting estimators.

Results: To complete objective (1), 41 articles were selected from the review process. Most of the included articles provide a solution related to parametric regression, and four communication schemes were identified among the presented methods. Most articles present a methodology that does not require any communication from the coordinating centre to data nodes. For objective (2), six approaches were extracted for conducting standard GLM-based statistical analysis. However, these approaches assumed evenly and identically distributed data across nodes. For objective (3), statistical procedures were derived to accommodate uneven node sample sizes and heterogeneous data distributions across nodes. Detailed algorithms were developed to highlight information-sharing requirements and operational complexity. Two approaches did not require any exchanged quantities from the coordinating centre to data nodes, while others needed sharing aggregated estimates and/or average of local gradients and Hessians.

Conclusions: This paper contributes to the field of healthcare research by providing an overview of the methods that can be used with horizontally partitioned data, by adapting these methods to the context of heterogeneous health data and by clarifying the workflows and quantities exchanged by the methods discussed. Documenting these adapted methods contributes to making the field of distributed data analytics more accessible to data custodians and health researchers, especially in the context of inference, which is not the focus of most of the existing HPSA literature. Further analysis of the confidentiality preserved by these methods

is needed to fully understand the risk associated with the sharing of summary statistics.

(JMIR Preprints 12/10/2023:53622)

DOI: <https://doi.org/10.2196/preprints.53622>

Preprint Settings

1) Would you like to publish your submitted manuscript as preprint?

Please make my preprint PDF available to anyone at any time (recommended).

Please make my preprint PDF available only to logged-in users; I understand that my title and abstract will remain visible to all users.

Only make the preprint title and abstract visible.

✓ **No, I do not wish to publish my submitted manuscript as a preprint.**

2) If accepted for publication in a JMIR journal, would you like the PDF to be visible to the public?

✓ **Yes, please make my accepted manuscript PDF available to anyone at any time (Recommended).**

Yes, but please make my accepted manuscript PDF available only to logged-in users; I understand that the title and abstract will remain visible.

Yes, but only make the title and abstract visible (see Important note, above). I understand that if I later pay to participate in <http://www.jmir.org/>

Original Manuscript

Original Paper

Félix Camirand Lemyre^{1,2,†}, Simon Lévesque^{1,2,4,†}, Marie-Pier Domingue^{1,2,5},
Klaus Herrmann², and Jean-François Ethier^{1,3,4,*}

1 GRIIS, Université de Sherbrooke, <https://griis.ca/>

2 Département de mathématiques, Faculté des sciences, Université de Sherbrooke

3 Département de médecine, Faculté de médecine et des sciences de la santé, Université de Sherbrooke

4 Health Data Research Network Canada

5 Chaire MEIE Québec - Le numérique au service des systèmes de santé apprenants

* Corresponding author

† These authors contributed equally to this work.

* Corresponding author

Distributed Statistical Analyses: A Scoping Review and Examples of Operational Frameworks Adapted to Health Analytics

Abstract

Background: Data from multiple organizations are crucial for advancing learning health systems. However, ethical, legal, and social concerns may restrict the use of standard statistical methods that rely on pooling data. Although distributed algorithms offer alternatives, they may not always be suitable for health frameworks.

Objective: This paper aims to support researchers and data custodians in three ways: (1) providing a concise overview of the literature on statistical inference methods for horizontally partitioned data; (2) describing the methods applicable to generalized linear models (GLM) and assessing their underlying distributional assumptions; (3) adapting existing methods to make them fully usable in health settings.

Methods: A scoping review methodology was employed for the literature mapping, from which methods presenting a methodological framework for GLM analyses with horizontally partitioned data were identified and assessed from the perspective of applicability in health settings. Statistical theory was used to adapt methods and to derive the properties of the resulting estimators.

Results: From the review, 41 articles were selected, and six approaches were extracted for conducting standard GLM-based statistical analysis. However, these approaches assumed evenly and identically distributed data across nodes. Consequently, statistical procedures were derived to accommodate uneven node sample sizes and heterogeneous data distributions across nodes. Workflows and detailed algorithms were developed to highlight information-sharing requirements and operational complexity.

Conclusion: This paper contributes to the field of health analytics by providing an overview of the methods that can be used with horizontally partitioned data, by adapting these methods to the context of heterogeneous health data and by clarifying the workflows and quantities exchanged by the methods discussed. Further analysis of the confidentiality preserved by these methods is needed to fully understand the risk associated with the sharing of summary statistics.

Keywords: Distributed algorithms; Generalized linear models; Horizontally partitioned data; Learning health systems

Introduction

Health Analytics at Scale

Learning health systems (LHS) are coming of age and are being deployed to address important health challenges at different scales. The framework starts by leveraging health data created across various activities. It obviously includes data points from clinics and hospitals, but the perimeter of

data required to meaningfully and optimally address important problems is much wider and includes research cohorts, biobanks, quantified self data, environmental exposures and social service delivery. While some questions might be addressed at the scale of an individual organisation, LHS focus on systems interactions and often require the analysis of processes and outcomes from various organisations. For example, to fully understand a cancer care trajectory, multiple data sources from multiple organisations will need to be examined to cover all relevant aspects (both within the traditional health system, but also in the community). This often implies at least regional organisations or bigger (provinces, states, countries), like in the context of the Health Data Research Network Canada (HDRN) or Health Data Research UK. Similarly, comparing various approaches is often a fruitful way to identify the best approaches and understand what works, why, and how to scale the promising projects. It can also be a way to amass a critical number of observations in the context of rarer diseases for example. Nevertheless, working with data from multiple data sources, from multiple organisations and located in multiple jurisdictions poses significant challenges.

Traditionally, the analytical methods used by researchers in health-related domains and others have relied on data pooling (sometimes referred to as data centralisation): all required data is physically copied to a single location where analysis can take place. However, when working with data from multiple jurisdictions (even when part of the same country like the Canadian provinces and territories), data pooling is often very difficult if not impossible for ethical, legal and social acceptability reasons.

There is therefore a pressing need to offer analytical methods allowing the analysis of such data without requiring the need to physically copy the data in a central location.

The primary intent of this paper is to lay the foundations for future practical assessments of the feasibility of conducting distributed statistical analyses in health-related contexts. We achieve this by reviewing the literature on existing methods, evaluating which ones can be applied to a class of widely used regression models, and precisely identifying their operational and information-sharing requirements in a notational framework that allows for straight comparisons. This unified framework enables the description of methods' operational workflows, quantities exchanged, and algorithmic implementation, and operates under assumptions commonly satisfied in health analytics.

The paper is structured as follows. We begin with a formal description of the distributed analytical framework considered in this study, followed by a discussion on the challenges associated with its implementation in health analytics. Next, we state our specific objectives and outline the methodology used to achieve them, including a scoping review and our approach to establishing the common notational framework for method description and comparison. After presenting the results pertaining to each objective, we discuss our findings and remaining challenges regarding distributed statistical analyses for health analytics.

Distributed Analysis

This paper is concerned with frameworks where the data needed for a statistical analysis consists of the data about n individuals (referred to as the *analytical dataset*), which are not all stored in a single source but are partitioned among K locations which will be called *nodes* hereafter. The mereological sum of all the data held at each node therefore forms the analytical dataset. Data can be partitioned horizontally or vertically (or in a mixed way).

- A horizontal partition implies that all data pertaining to a given individual can be found in a single node. If we assume that patients receive care only in one province, Canadian provincial health administrative datasets hosted by organizations like Population Data BC, ICES in Ontario or the Manitoba Centre for Health Policy (MCHP) in Manitoba can be part of a horizontal partition. A clinical trial where each recruiting site captures all data for a given subject is another example.
- A vertical partition occurs when all data of a certain type is available in a single node for a group of individuals. A classic example is a hospital with its various information systems. All

pathology results can be found in the pathology system, all billing information can be extracted from the finance system, all X-rays are accessible in the picture archiving and communication system (PACS), etc. But to get the full picture of the care received by a patient, multiple systems need to be interrogated. Similarly in the research setting, health administrative data may be in a provincial data centre and genomics data could be held in a research institute.

A mixed partition occurs when both principles partly apply: some individuals may have their data spread out across nodes, and different individuals may be present in different nodes.

Assumptions

The difficulties in conducting analyses on a large scale mentioned above are often associated with horizontally partitioned data, and the current work focuses on this type of partition. The methods presented in this article might therefore not be directly applicable to vertically partitioned data.

One group of approaches often labelled as *distributed analysis* involves calculations at each participating node and exchanges of the resulting aggregated statistics with a *coordinating centre* (CC), which can itself also perform additional calculations based on the received aggregated statistics. The CC can be an organisation not responsible for a data node or a data node taking the additional role of CC for a given analysis.

It is important to note that whether in the more traditional way of data pooling or using distributed approaches (where the data is not copied centrally), data sources will be different on multiple levels. They will represent information using data models with significant variability in terms of structure and technology, but also in terms of semantics. This situation also leads to heterogeneous data where the presence of predictors and outcomes is likely to be different in different nodes. Different approaches (e.g. data mediation or extract-transform-load) have been developed to address these issues, and the current work assumes that one of them has been applied so that the data nodes mentioned hereafter are assumed to share the same structure, the same technological syntax and the same semantics as well as no missing data.

Horizontally Partitioned Statistical Analytics

In what follows, the field that pertains to the statistical analysis of horizontally partitioned and semantically homogeneous data that cannot be consolidated into a central location will be called *Horizontally Partitioned Statistical Analytics* (HPSA).

Methodological contributions to this field have arisen from several streams of literature. Meta-analysis and meta-regression methods (see e.g. [1]) can be viewed as part of HPSA, e.g. by considering that each node-specific dataset belongs to a different "study". However, their scope is narrower compared to HPSA because they typically assume that only established study-level estimates are available as data. Conversely, HPSA allows for the sharing of additional summary statistics between the nodes and the CC, such as gradients and Hessians to ensure the best possible performance at the global level. Since meta-analysis does not leverage any supplementary information that could be obtained from studies with access to patient-level data, it can be susceptible to biased estimation, especially in settings with rare outcomes or in the presence of data nodes with limited sample sizes [2]. As meta-analysis and meta-regression methods have been extensively covered in the literature, approaches specifically designed for the analysis of already-established study-level estimates will not be discussed hereafter.

An important research community that has generated a significant amount of analytical contributions is concerned with the massive data setting. There, a dataset often cannot be processed by a single server and is therefore split across multiple machines, which are then considered as nodes able to perform computations and send aggregated results to a CC tasked with fitting a global model from them. The methodological avenues proposed in this setting share similarities with the ones designed for the multi-research facility setting involved in LHS, but also have important differences. For

example, in the massive data setting, the experimenter has control over the distribution of individuals across nodes, which is typically not the case in multi-research facility studies. So while these approaches share mechanistic similarities and have been suggested as options to consider in the health domain, some hypotheses may not hold. In regression settings, it is often reasonable to assume that the regression link between the response and covariate predictors is the same across nodes. However, assuming that the sampling distribution of covariates involved is equal across nodes is unrealistic in the health domain, particularly due to the presence of data centres that may systematically involve different types of patients. For example, certain clinics may predominantly serve older individuals. While this may not affect the estimation of parameter values, it can have implications for computing confidence intervals to ensure the validity of inferences.

So far, two reviews discussing methods applicable to horizontally partitioned data have been published in the literature [3][4]. However, their focus is on the massive data setting, which works almost invariably under the assumption of even sampling distribution of covariates and equal sample sizes across nodes, and statistical inference tasks beyond parameter estimation are barely covered. This makes them less helpful for health analytics purposes since most studies involving data analyses rely on confidence intervals or hypothesis testing in settings where predictors' distribution and sample sizes vary across nodes.

Contemporary challenges in HPSA

The problem is threefold. First, there is a need to raise awareness regarding the existence of HPSA approaches among researchers aiming at undertaking statistical analyses from horizontally partitioned data, especially in health analytics. The reflex is often to request data pooling because it is perceived as the sole option. This has been the tendency of requests made by researchers to HDRN Canada. Practitioners are usually concerned with finding the most appropriate statistical model that will take into account as many of the features of their specific context of application as possible. Consequently, a clear and unifying mapping of the state of the HPSA field is needed for them to be informed of the scope of existing methods available for their analyses to see whether alternatives to pooling exist.

Second, as underlined above, methodological contributions came from research fields whose working assumptions can be fundamentally different from the ones researchers would be willing to assume in health analytics. To ensure proper use of statistical inference techniques, it is necessary that the underlying assumptions of existing methods be adequately identified and understood. If necessary, these methods should be adapted to suit the specific requirements of health applications, thereby ensuring accurate and reliable results.

Third, data custodians have to be properly informed on data-sharing requirements entailed by the use of a specific HPSA method applicable to a given research setting. While HPSA avoids the complexities of pooling data, there are still flows of information that have to be acceptable to data stewards. However, even in basic statistical scenarios, available methods are often presented in a way that makes them challenging to compare in terms of information-sharing requirements and operational complexity. Therefore, there is a need for clearer and more accessible presentations of these methods to facilitate decision-making regarding data sharing and operational implementation.

Although it would be ideal to offer managers a comprehensive operational workflow for each identified method to evaluate the information shared and execution complexity, with their accompanying underlying modelling assumptions, the abundance and diversity of available approaches make it unfeasible to accomplish this in a single paper. In fact, methods often differ in terms of their targeted application beyond their distributed aspect. For example, differences may exist in the studied model (e.g., linear, logistic or Cox regression, additive models), the dimensionality/sparsity of the predictor variable space, use of regularization or shrinkage, the presence of missingness, confounders, imbalances, heterogeneity, etc.

Objectives

The objectives of this article are:

1. To identify and map, from the literature, methodological approaches that make it possible to perform confidence intervals estimation and hypothesis testings from a horizontally partitioned dataset;
2. Among the approaches identified, to describe the ones that allow to conduct generalized linear model (GLM) analyses, and to identify their distributional assumptions;
3. Based on the approaches identified for GLM-based inferences, to present methods adapted to the setting of uneven sampling distributions across nodes, and to compare them in terms of information-sharing requirements and operational complexity.

A scoping review methodology was chosen to achieve objective O1 of mapping the state of the field of HPSA that pertains to inference procedures. For our second objective (O2), we identified, from the articles selected from the literature search, the ones that presented a methodological framework for conducting statistical inference procedures from a GLM with horizontally partitioned data. We then used these frameworks to derive and describe GLM estimators that are applicable to horizontally partitioned datasets. For each identified method, we analyzed and reported its communication workflow and the distributional assumptions. For our third objective (O3), we first used statistical theory to adapt the identified procedures to the unequal sample size and uneven covariate distribution setting. Algorithms and mathematical expressions for the quantities involved are reported. For conciseness, we present mathematical formulas for estimation procedures of confidence intervals only. Expressions involved for hypothesis testing are similar and can be deduced following the close connecting between confidence intervals and hypothesis tests in GLMs, see e.g. [5].

The mathematical description of the GLM setting considered for this analysis is described below, along with mathematical notations to be used.

Mathematical foundations and notation

In the following, lowercase bold letters will represent vector-valued quantities, while uppercase bold letters will denote matrices. The j^{th} element of any vector $a \in \mathbb{R}^p$ will be denoted as $[a]_j$. Similarly, the entry at position (j, l) of any matrix $A \in \mathbb{R}^{p \times p}$ will be denoted as $[A]_{jl}$. If g is a real-valued and invertible function, we will use $g^{(-1)}$ to represent its inverse. Additionally, if f_θ is a real-valued function that depends on a parameter vector θ and is twice continuously differentiable, $\nabla_\theta f_\theta$ and $\nabla_\theta^2 f_\theta$ will respectively indicate the gradient and Hessian matrix of f_θ with respect to θ .

Model mathematical assumptions

A mathematical depiction of the horizontally partitioned data framework studied in this paper is as follows. There are n individuals horizontally partitioned across K data storage nodes. Each node's dataset is denoted by $D^{(k)} = \{z_i^{(k)} = (x_{1i}^{(k)}, \dots, x_{pi}^{(k)}, y_i^{(k)})^T\}_{i=1}^{n^{(k)}}$, where $1 \leq k \leq K$. Here, $z_i^{(k)}$ represents the measurements on the i^{th} individual at node k , where $y_i^{(k)} \in \mathbb{R}$ denotes their response variable and $[x_{1i}^{(k)}, \dots, x_{pi}^{(k)}]^T \in \mathbb{R}^p$ denotes their covariate vector. The total sample size at node k is denoted by $n^{(k)}$. The combined datasets $D^{(1)}, \dots, D^{(K)}$ make up the whole dataset without any duplicated individuals, indicating that $\sum_{k=1}^K n^{(k)} = n$.

Throughout the analysis, it is assumed that the $z_i^{(k)}$'s are independent across $1 \leq i \leq n^{(k)}$ and $1 \leq k \leq K$, and there is no missing data. Additionally, the size of the covariate space (i.e., the dimension of $[x_{1i}^{(k)}, \dots, x_{pi}^{(k)}]^T$, which is equal to p representing the number of features to include as predictors in the GLM) is assumed to be low, eliminating the need for regularization or variable selection. Finally, it is

assumed that each node possesses a non-negligible proportion of the whole dataset. Specifically, for each $k \in \{1, \dots, K\}$, the quantity $n^{(k)}/n$ is bounded away from 0 and 1 as the sample size n tends to infinity, denoted as $n^{(k)}/n \rightarrow p^{(k)} \in (0, 1)$.

Mathematical description of the GLM framework

The formulation of the GLM considered in this article encompasses various commonly used regression models such as linear regression, logistic regression, Poisson regression, and probit models. It assumes that the density or probability mass function of each response variable (known as the random components) belongs to the exponential family of distributions. Within this formulation, the mean of the response variable is expressed as a function of a linear combination of the corresponding covariate vector. Formally, it assumes that there exist unknown parameters $\beta^* \in R^{p+1}$ and $\phi^* > 0$, and known model-specific functions b, c, g, h such that with $x_i^{(k)} = [x_{0i}^{(k)}, x_{1i}^{(k)}, \dots, x_{pi}^{(k)}]^T$ and $x_{0i}^{(k)} = 1$,

$$y_i^{(k)} | x_i^{(k)} \sim f(\cdot; x_i^{(k)}, \beta^*, \phi^*),$$

where for any $\beta = [\beta_0, \beta_1, \dots, \beta_p]^T \in R^{p+1}$ and ϕ ,

$$f(y; x_i^{(k)}, \beta, \phi) = \exp \left[\frac{y h(\beta^T x_i^{(k)}) - b(h(\beta^T x_i^{(k)}))}{\phi} + c(y, \phi) \right]. \quad (1)$$

In formula (1), b is such that $b'(h(\beta^T x_i^{(k)})) = E(y_i^{(k)} | x_i^{(k)}) = g^{-1}(\beta^T x_i^{(k)})$, with $b'(x) = \partial b(x) / \partial x$. In this framework, g is called *link function*, the term $h(\beta^T x_i^{(k)})$ is usually referred to as the *natural parameter* and b as the *cumulant function*. ϕ is often called the *dispersion parameter*, and is either known (e.g. with $\phi=1$) or unknown. When $h(x)=x$ (i.e. h is the identity function), the link g is called *canonical*. The logistic regression model is obtained upon taking $\phi=1$, $h(x)=x$, $b(x)=\log(1+e^x)$, $c(y, \phi)=0$ and $g(x)=\log\{x/(1-x)\}$. The linear regression model with homoskedastic residual error variance ϕ is derived upon setting $h(x)=x$, $b(x)=x^2/2$, $c(y, \phi)=-y^2/(2\phi)-\log(2\pi\phi)/2$ and $g(x)=x$. Hence, both the logistic and the linear regression models rely on a canonical link function in the exponential family distribution.

Methods

Methodology related to objective O1

Scoping reviews are well-suited to efficiently map key concepts within a research area [6]. They are widely acknowledged for their ability to clarify working definitions and conceptual boundaries in a specific topic or field [7], facilitating a shared understanding among researchers regarding the status of the research area. These considerations make the scoping review methodology well-designed to achieve objective O1.

Scoping studies utilize systematic searches of relevant databases, employing specific keywords to define the boundaries of the research field. However, identifying these keywords can be challenging, particularly when relevant papers are scattered across different research streams or in independent clusters that do not reference each other. To address the risk of overlooking significant methodological contributions due to a limited number of keywords, a snowballing literature search was initially conducted to generate a comprehensive list of keywords related to HPSA. The scoping review then proceeded with a systematic literature search using the identified keywords. It's worth noting that, since the planning of the scoping study is independent of the search approach, the guidelines presented in [6] are still appropriate.

Methodology pertaining to the snowballing keywords search

Snowballing is generally used as a literature search method aiming at identifying papers belonging to a given field [8]. It typically consists in three steps:

1. Initiate searches in prominent journals and/or conference proceedings to gather an initial set of papers.
2. Conduct a backward review by examining the reference lists of the relevant articles discovered in steps 1 and 2 (continue iterating until no new papers are found).
3. Perform a forward search by identifying articles that cite the papers identified in the previous steps.

To avoid selection bias, the initial set of papers for the snowballing approach in (1) is sometimes generated through a search in Google Scholar (see e.g. [9]). The latter strategy was used here too.

As mentioned earlier, here, the snowballing search strategy was used in preparation for the application of the scoping review protocol, with the goal of identifying relevant keywords. Specifically, the starting set of papers was assembled by screening titles and abstracts from the first 50 papers generated through a Google Scholar search using the strings *distributed inference* and *federated inference*. The main inclusion criterion was "presents, applies or discusses a statistical inference method to analyse horizontally partitioned data". Then, the backward and forward snowballing steps approaches were applied.

From the set of keywords found in the selected papers, a list of those relevant to HPSA but not directly associated with any specific method was retained for the scoping review step. It is worth noting that, since the objective of the scoping review is to identify statistical inference methods for horizontally partitioned data, keywords linked to method identifiers have to be excluded from the retained list to avoid pre-selection bias in the scoping review phase of this project.

Selected keywords that were identified from the snowballing literature search are *distributed algorithms*, *distributed estimation*, *distributed inference*, *distributed learning*, *distributed regression*, *federated inference*, *federated estimation*, *federated learning*, *privacy-protecting algorithm*, *privacy-preserving algorithm* and *aggregated inference*.

Methodology pertaining to the scoping review

The scoping review's methodological framework of Levac et al. [10] (see also [6]) was followed. The steps are briefly described below. A detailed protocol is available in Appendix A.

We conducted a comprehensive search across four bibliographic databases, namely (1) MEDLINE, (2) Scopus, (3) MathSciNet, and (4) zbMATH, to encompass the interdisciplinary nature of the topic and identify relevant research articles. Our search strategies were based on two key concepts: distributed data and statistical inference. In addition to the keywords obtained from the snowballing step, we incorporated terms like *confidence interval* to target articles focusing specifically on statistical inference. To ensure the inclusion of recent advancements, our search was limited to papers published from 2000 onwards. This cutoff date was chosen to account for the emergence of distributed data, the prevalence of massive datasets, and advancements in technology. It was set conservatively to capture any early-developed methods and ensure comprehensive coverage of the topic.

After completing the primary research, a two-stage selection process was employed. Initially, two authors (MPD, FCL) collaborated to screen all articles identified through the search strategy based on their titles and abstracts. Subsequently, the full texts of the selected articles were independently reviewed by both authors to finalize the selection process. This rigorous approach ensured a thorough evaluation of each article's relevance and eligibility for inclusion.

The primary inclusion criterion for the selection process was as follows: *Presents a solution for conducting inferential statistics on horizontally partitioned data*. This criterion was utilized to ensure that the chosen articles specifically addressed the methods associated with performing statistical

inference on horizontally partitioned data.

The following exclusion criteria were derived directly from objective O1:

1. Does not address inferential statistics, including confidence intervals, hypothesis testing, or asymptotic normality.
2. Does not provide a methodological contribution.
3. Presents a solution for encryption or secret-sharing.

To ensure the inclusion of validated approaches, the selection process only considered published papers that had full-text availability in English or French. Discussion papers were excluded as they do not present novel methods or approaches.

Exclusion was considered if any of the exclusion criteria were met or if any of the inclusion criteria were not met.

Finally, the references of each included article from the databases were assessed to identify any relevant articles that may not have been captured during the initial screening due to specific keywords. This additional step in the selection process was necessary given the broad range of vocabulary used to describe applicable approaches in our context.

Data extraction for the included articles was conducted by one author (MPD) and followed a collectively developed data-charting form. Model type (*parametric regression*, *semi-parametric regression*, *non-parametric regression* or *not specific to regression*) and number of communications from CC to nodes (0 or ≥ 1) were among the data extracted. All methods from the included articles were subsequently classified according to their specified characteristics, as outlined in the protocol. Additionally, as part of the analysis, we conducted a screening of the general distributed approaches commonly employed across all specific methods.

Methodology related to objective O2

To achieve objective O2, three steps were taken. First, we identified methodological approaches from articles included in the scoping review that enable parameter and confidence interval estimations from horizontally partitioned data within a standard GLM framework. Methods designed specifically for the particular cases of linear or logistic regression were also reported but were not analyzed in detail. Second, we extracted workflows for each approach to determine the information exchanged between data storage nodes and the CC. Third, we analyzed the mathematical assumptions necessary for parameter estimation and the consistency of confidence interval procedures. We specifically reported the assumptions related to the distribution of node-specific covariates.

Identification of the approaches

To identify approaches that enabled the fitting of any GLM using horizontally partitioned data, two authors (FCL, MPD) independently assessed all articles included in the scoping study. The reviewers specifically looked for articles that discussed approaches applicable to the GLM class described in the section Mathematical [foundations and notation](#), including likelihood-based methods, M-estimation, and estimating equations. Additionally, we identified and reported articles that specifically focused on regression settings for linear or logistic regression. However, unless the method described was considered easily adaptable to the GLM framework, these articles were not retained for detailed analysis.

A method was selected if it provided an algorithm for fitting GLMs using horizontally partitioned data, aligning with the characteristics outlined in the section Mathematical [foundations and notation](#). In cases where an article presented asymptotic normality results for the estimators but did not provide an estimator for the asymptotic variance-covariance matrix, the article was still retained, and an estimator for the asymptotic variance was derived using the available calculated quantities.

Since our GLM framework assumes no missing values, low dimensionality, and a small number of nodes relative to the total sample size, any terms related to these specific conditions mentioned in an article's methodology were disregarded. Consequently, the calculations for confidence intervals were

adjusted accordingly. If an article solely focused on one of these aspects without contributing to the overall methodology, it was not included in the final selection.

Methodological components regarding parameter estimation and confidence interval procedures were extracted from the screened articles. Specifically, the focus was on understanding how parameters should be estimated within a horizontally partitioned framework and how confidence intervals should be computed for these parameters. For each article, the formulas related to quantities shared among the nodes and quantities calculated by the CC were derived and analyzed. These formulas were examined within a workflow that indicated the necessary circulation of information for the procedure to be executed. The derived workflows constituted the first part of our defined unified framework for HPSA approach comparisons in GLM settings.

For the reported results, the rationale behind each method that was deemed suitable for fitting GLMs was documented, along with the corresponding reference to the paper included in the scoping study where the method was introduced or discussed.

Articles that discussed approaches specifically applicable to the cases of linear or logistic regression were also mentioned, but not elaborated on in detail.

Methodology related to objective O3

In most statistical settings with horizontally partitioned data, it is commonly assumed that the sample sizes of the data nodes are equal, and that the distribution of covariates is the same across all nodes. However, when the number of nodes is fixed and relatively small compared to the sample sizes, it is possible to adapt a particular approach to handle situations where the sample sizes and covariate distributions vary across nodes. This can be achieved by combining the theoretical arguments presented in the original article of the method with the principles of asymptotic statistics theory concerning maximum likelihood estimation.

To adapt a given approach for situations where sample sizes and covariate distributions differ across nodes, the following steps were taken:

1. The formulas for the relevant quantities were modified to emphasize the changes caused by this scenario. It was ensured that the adapted quantities were equivalent to their counterparts presented in the original article for an equal sample size setting.
2. Using asymptotic theory, an asymptotic normality result was derived for the estimators of interest, assuming a set of assumptions that accommodated potential variations in sample sizes and covariates sampling distribution across nodes, while still enabling meaningful theoretical arguments.
3. Formulas for the asymptotic variances were derived. Statistical theory on maximum likelihood estimation was employed to obtain consistent estimators for asymptotic variances. The latter estimators were derived under the constraint that they had to be calculated without requiring any additional communication round between the CC and the nodes. Thus, throughout the adaptation process, the communication workflow remained unchanged compared to the original method.

These steps ensured the mathematical correctness of adapting the approaches to handle different sample sizes and covariate distributions across nodes. Importantly, these adaptations maintained consistency with the original method's communication workflows. To adapt and compare the methods based on their information-sharing requirements, common assumptions and a unified mathematical and algorithmic notation were necessary. Ultimately, these assumptions and the notation, along with workflow types derived for O2, enabled us to establish a unified notational framework for approach comparisons when performing HPSA based on GLMs in the context of health analytics.

We describe the statistical estimates of interest. A standard GLM typically includes one or two

unknown parametric components. The first are the β parameters, which are commonly assumed to be unknown. The second parameter is the nuisance parameter ϕ , which can either be known (e.g., in logistic models) or unknown (e.g., in linear models). In practical applications, when ϕ is unknown, its estimated value is often not the main focus, although the latter is necessary to estimate the asymptotic variance of the β parameter estimates.

In the upcoming analysis, we will assume that the parameter ϕ is unknown and estimated using the recommended approach in the selected methods. However, in the case where ϕ is known, the process becomes simpler. This involves substituting the known value of ϕ and disregarding the estimation step. It is important to highlight that estimating ϕ requires additional information to be shared between the nodes and the CC, but it does not necessitate any extra communication round between them.

The estimation process for both the β parameters and the ϕ parameter are discussed. Additionally, we explained how to compute an estimator for the asymptotic variance specifically for the estimator of β^* . It is important to note that the results presented below can be modified and extended to develop a similar procedure for estimating ϕ^* .

Using these results, based on an estimator of β^* , say $\hat{\beta}$ and a formula for the estimator of the asymptotic variance-covariance matrix involved in its associated asymptotic normality result, say $\hat{\Sigma}$, Wald-type $(1-\alpha)$ confidence intervals can be computed for each component of β^* using the formula

$$[\hat{\beta}]_j \pm z_{1-\alpha/2} \sqrt{[\hat{\Sigma}]_{jj}/n} \text{ for } j \in \{0, \dots, p+1\}.$$

Regarding the reported results, for each approach considered, we presented the formulas necessary to compute the final estimates of the β parameters and their corresponding confidence intervals. The presentation of these formulas was designed to emphasize the communication workflow. Furthermore, a comprehensive algorithm was provided, outlining the step-by-step process.

In addition, the asymptotic normality of the β parameter estimators was stated, accompanied by the formula for the asymptotic variance and its consistent estimator. Detailed proofs for these results can be found in Appendix B.

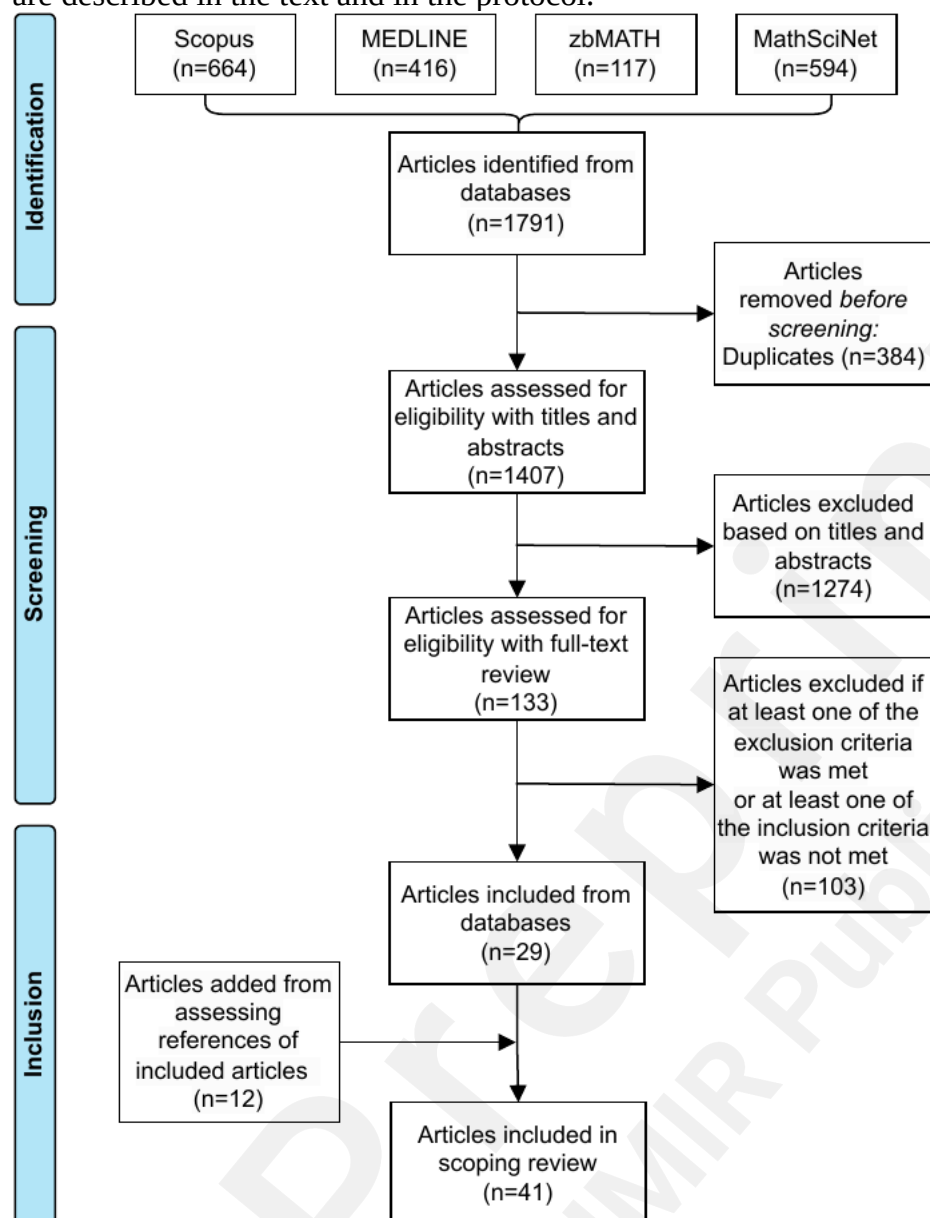
Results

Results related to objective O1

Search outcomes from the scoping review

As presented in Figure 1, a total of 1407 articles were initially identified across all four databases after removing duplicates. Subsequently, a majority of these articles ($n=1274$) were excluded based on the evaluation of titles and abstracts, leaving 133 articles for eligibility assessment through full-text review. Following this assessment, 29 articles were included from the databases. Additionally, by reviewing the references of the included articles, 12 more articles were identified and added to the study.

Figure 1. Article selection process for the scoping review. Detailed inclusion and exclusion criteria are described in the text and in the protocol.



Among the additional 12 articles obtained through the assessment of references from included articles, it is observed that most of them did not mention statistical inference or related terms in their abstracts (e.g., [2][11][12]). Consequently, these articles were not captured in the initial database search results. Furthermore, some articles directly referred to the specific method used without including any keywords related to horizontally partitioned data in their abstracts or titles (e.g., [13] [14]), which greatly reduced the chance of initially identifying them. However, during the process of reviewing the references of included articles, all the relevant papers that were initially identified through the snowballing strategy were eventually retrieved either through the search strategy or the selection process based on the references of included articles.

Results of the scoping review

Each article included in the scoping review put forth one or multiple methodological approaches pertaining to objective O1. Similarities and differences regarding the communication schemes involved and their background of origin are summarized below.

First, all selected articles discuss one or more statistical procedures that operate on horizontally partitioned data using one of the communication schemes depicted in Figures 2 to 5.

- In Workflow I, as shown in Figure 2, each node calculates summary statistics from its own samples, and the results are sent to the CC. The CC combines the information provided by each node to produce the final estimates. This communication approach is commonly referred to as a "one-shot" or "non-iterative" in the literature, although not always consistently.
- In Workflow II, as shown in Figure 3, multiple communication rounds are allowed between the CC and the data storage nodes. This allows for iterative interactions between the nodes and the CC to refine the estimates.
- Some approaches fundamentally differ from the two previous workflows by assigning a different role to one of the nodes, say node 1, compared to the others. These approaches operate using Workflow III as illustrated in Figure 4, where node 1 follows a distinct communication pattern compared to the other nodes. In the papers included in the scoping review that discuss these approaches, node 1 is invariably designated as the CC. However, in the context of the current paper, their roles were distinguished. The additional step performed by the CC, which involves data aggregation, can be particularly well-suited for privacy protection purposes in practice.
- The particular setting shown in Workflow IV in Figure 5 requires two back-and-forth communication exchanges between each node and the CC at each iteration. This communication pattern distinguishes it from the other workflows.

In light of the preceding discussion, from an operational standpoint, two categories of workflows emerge. On one hand, there are workflows that do not necessitate any communication from the CC to the nodes, which are captured in Workflow I. On the other hand, there are workflows that involve one or more communication exchanges from the CC to the nodes, which are captured in Workflows II, III and IV.

Figure 2. Workflow I: each node calculates summary statistics from its own samples. Results are sent to the CC which combines the information provided by each node to produce the final estimates.




Figure 3. Workflow II: multiple communication rounds are allowed between the CC and the data storage nodes.




Figure 4. Workflow III: multiple communication rounds are allowed between the CC and the data storage nodes, with node 1 following a distinct communication pattern compared to the other nodes.

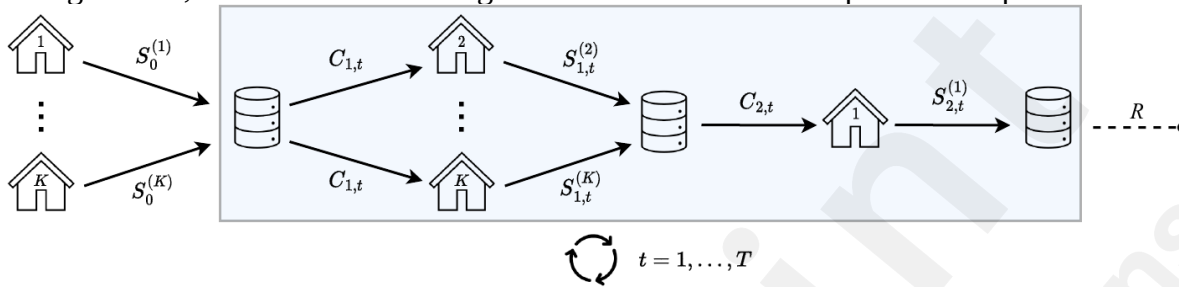
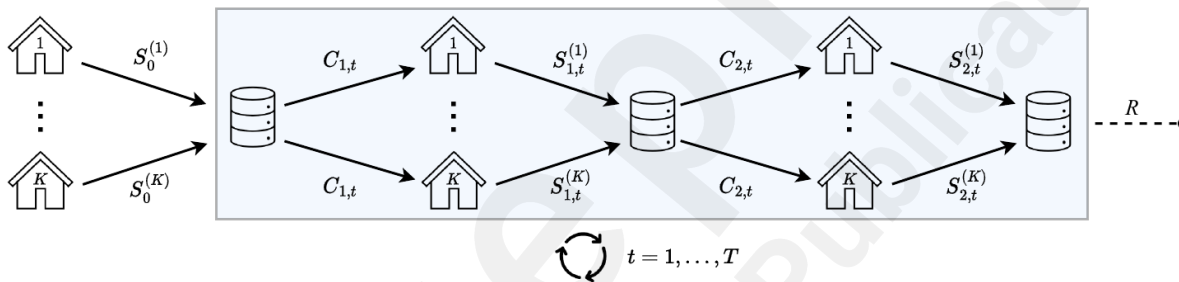


Figure 5. Workflow IV: multiple communication rounds are allowed between the CC and the data storage nodes, with two back-and-forth distinct communication exchanges between each node and the CC at each iteration.



In order to emphasize similarities among the methods presented in these articles and facilitate the identification of methods suitable for specific purposes, a systematic classification is presented in Table 1. The articles are categorized based on the type of models employed and the number of communications from the CC to the individual nodes.



Table 1. Classification of articles included in the scoping review.

Type of model	0 communication from CC to nodes	≥ 1 communication from CC to nodes
Parametric regression	Basiri, Ollila and Koivunen [15]; Battey et al. [16]; Fan, Guo and Wang [17]; Guo, Sun and Jiang [18]; Chen and Xie [14]; Lin and Xi [19]; Rosenblatt and Nadler [12]; Zhang, Duchi and Wainwright [20]; Chang, Bu and Long [21]; Wu et al. [22]; Hector and Song [23]	Huang and Huo [24]; Jordan, Lee and Yang [25]; Mozafari-Majd and Koivunen [26][27]; Yue, Kontar and Gómez [28]; Duan, Ning and Chen [29]; Duan et al. [2]; Tong et al. [30]; Di, Wang and Lian [31]; Edmondson et al. [32]; Luo and Li [33]; Shu, Young and Toh [34]
Semi-parametric regression	Zhao, Cheng and Liu [35]; Park et al. [36]	Luo, Sun and Zhou [37]; Duan et al. [11]
Non-parametric regression	Liu, Shang et Cheng [38]; Zhang et al. [39]; Volgushev, Chao and Cheng [40]	Wang et al. [41]
Not specific to regression	Atta-Asiamah and Yuan [42]; Minsker [43]; Lin and Xi [44]; Bruce et al. [45]; Chen and Peng [46]; Nezakati and Pircalabelu [47]; Banerjee, Durot and Sen [13]; Shi, Lu and Song [48]; Wu et al. [49]	Lai, Hanning and Lee [50]

The majority of the methods were published within the methodological setting of Big or Massive data/Multi-machine, while some were reported within the context of health analytics. Within the Big or Massive data/Multi-machine methodological setting, many methods involve an initial step of random data partitioning among multiple machines. However, certain methods assume a scenario where data is already stored on separate machines, as observed in [17] and [25]. Furthermore, it is worth noting that no articles published prior to 2010 were included, aligning with our initial hypothesis regarding the identification of contemporary methodological settings. The majority of the included articles (30 out of 41) were published after the year 2018.

The majority of articles address a setting where a CC exists external to the nodes, as exemplified by articles such as [44], [40], and [28]. In contrast, as mentioned above, some articles designate one of the nodes to assume this central role, as demonstrated in [21].

The methods identified through our search strategy share a common characteristic of utilizing a global model that incorporates population-level parameters. In some cases, these parameters may also include node-specific components to accommodate node-specific statistical heterogeneity in the outcome-predictors relationship, which captures deviations from the population-level conditional probability distribution of the outcome given the predictors.

A few of the reported methods have the capability to yield results identical to those obtained if the individual line data were pooled from all nodes, see e.g., [22] and [34].

While most articles featured methods related to regression models, including semi-parametric and non-parametric designs, a few reported results for other modeling frameworks. These include methods for M-estimation, U-statistics, symmetric statistics, and natural parameter estimation, some of which encompass regression models as a specific instance.

Results related to objective O2

Six approaches were selected as applicable to the standard GLM framework. They all assumed that nodes had equal sample sizes and identical distributions for the covariates.

Simple averaging

One of the simplest methods for horizontally partitioned data analysis, often referred to as the "simple averaging method" or the "divide-and-conquer" approach, has been extensively studied in the literature, see [20] and [51] which were included in our scoping study. It operates through Workflow I in Figure 2. In this approach, node-level model estimates are gathered and averaged at the CC to generate the final estimates.

In the context of GLM, each node is initially tasked with calculating the maximum likelihood estimator (MLE) of the β^* and ϕ^* parameters using their respective data. Additionally, the Hessian matrix of the log-likelihood function with respect to the β parameters must be computed for constructing Wald-type confidence intervals. The estimated parameters and the computed Hessian matrix are then transmitted to the CC.

The final parameter estimates of β^* are obtained by averaging the node-specific estimates, while the local Hessians and estimates of ϕ^* are utilized to compute an estimator for the asymptotic variance.

Single distributed Newton-Raphson updating

The single distributed Newton-Raphson updating method is an iterative procedure that includes an additional communication round between the CC and the nodes, compared to the simple averaging method. It was originally proposed as the "distributed one-step" method in [24], but here it is referred to by a different term to avoid any confusion regarding communication complexity. This method operates using Workflow II, as depicted in Figure 3, with $T=1$ (where T represents the number of cycles in the iteration scheme). It enhances the simple averaging estimators by incorporating a single distributed Newton-Raphson updating step.

In the context of GLM, each node first calculates the MLE of β^* and ϕ^* , and transmits them to the CC. The CC aggregates these estimates using averaging and sends the result back to the nodes. The nodes then compute the gradient and the Hessian matrix of the log-likelihood function, evaluated at the received β^* and ϕ^* estimates. Subsequently, the gradient and the Hessian matrix are sent back to the CC, which averages them and computes a Newton-Raphson updating step based on the simple averaging estimates. An estimator for the asymptotic variance can be calculated by utilizing the received Hessian matrices and the updated estimate of ϕ^* .

Multiple distributed Newton-Raphson updatings

The multiple distributed Newton-Raphson updating method leverages the fact that, for standard GLMs, the algorithm typically used to calculate the MLE of β^* and ϕ^* in a centralized pooled setting can be executed in a distributed manner without any loss of information. This is possible because the algorithm relies on Newton-Raphson updatings (or sometimes Fisher scoring updatings) that are expressed using two sums of node-specific summary statistics, namely local gradients and local Hessian matrices of the log-likelihood function, evaluated at the β^* and ϕ^* estimates from the previous iteration. A version of this method is proposed in [22] under the logistic regression framework. It operates through Workflow II in Figure 3 for a general $T \geq 1$.

Distributed estimating equation

The class of estimating equations estimators is vast and encompasses a broad range of statistical estimation techniques, including likelihood-based approaches that rely on searching for critical points. The fundamental idea behind estimating equations methods is to establish a system of equations that involve both the sample data and the unknown model parameters. These equations are then solved to determine the parameter estimates. MLEs, which are obtained by setting the gradient of the log-likelihood function with respect to the unknown parameters equal to zero, belong to the class of estimating equations estimators.

The distributed estimating equations approach involves gathering summary statistics from nodes at the CC level, enabling the reconstruction of the estimating equations, or more commonly, an approximation of them that would have been obtained in a pooled centralized setting. This method is discussed in [19] and operates using Workflow I, as depicted in Figure 2.

In the context of GLMs, the distributed estimating equations approach involves initially assigning each node the task of computing and sending their local MLEs and the Hessian matrix of their local log-likelihood, evaluated at those MLEs, to the CC. The CC utilizes these received quantities to reconstruct the global estimating equations or an approximation thereof. This reconstruction ultimately leads to an analytical solution for obtaining the resulting estimates. Confidence intervals are computed using a combination of the Hessian matrices and the final estimator of ϕ^* .

It is important to note that when this approach is applied in the context of linear regression, it enables the acquisition of β^* parameter estimates that are identical to those obtained in a pooled centralized setting.

Distributed estimation using a single gradient-enhanced log-likelihood

This method differs fundamentally from the ones discussed thus far, as it involves a distinct role for one particular node in obtaining model parameter estimates. It operates using Workflow III, as depicted in Figure 4, and was proposed in [25] under the name "Surrogate likelihood". This approach relies on an approximation of the global likelihood by viewing it as an analytic function. It expands the global likelihood into an infinite series around an initial guess $\hat{\beta}_{SGE,0}$ and replaces the higher-order derivatives (order ≥ 2) of the global likelihood with those of a Taylor expansion of a node's (e.g., node $k=1$) local likelihood around the same value. By following this procedure, the so-called surrogate likelihood can be solved using data from node $k=1$ and aggregated gradients from nodes

$k \in \{2, \dots, K\}$.

In the context of GLM, the CC first collects the necessary information to compute initial estimates for the parameters β^* and ϕ^* . These initial estimates can be obtained through various approaches, such as a simple averaging estimator or the MLEs computed using data from node 1. These initial estimates are then transmitted to nodes $k \in \{2, \dots, K\}$. Each of these nodes calculates the gradient of the log-likelihood function, evaluated at the received estimates, and sends it back to the CC. The CC averages these gradients and sends the result to node 1. Node 1 solves a gradient-enhanced log-likelihood using its own data and the received average gradient. The resulting estimate is sent back to the CC as the final estimate. To compute confidence intervals, each node must send the Hessian matrix of its local log-likelihood function, evaluated at the initial received estimate.

The steps related to estimation can be repeated multiple times.

Distributed estimation using multiple gradient-enhanced log-likelihoods

This method is in the spirit of the *distributed estimation using a single gradient-enhanced log-likelihood* approach described above, except that all nodes have to solve a gradient-enhanced log-likelihood instead of only one of them. Results pertaining to statistical inference are discussed in [17] under a penalized setting. A non-penalized version of this method was introduced in [51] although the latter did not discuss confidence intervals or hypothesis testing, and hence was not included in our scoping review. It operates through Workflow IV depicted in Figure 5.

The following subsection presents the six approaches described in objective O2 within a unified notational framework, that accounts for the peculiarities commonly encountered in health analytics. Algorithms are derived using this common notation to rigorously describe the methods and enable their comparison. Theoretical results regarding the estimators involved are detailed and proven in Appendix B. While this section is necessary to increase trust in HPSA methods by transparently showing precisely what information is shared with the CC through comprehensive mathematical formulas, the summary provided in Table 2 suffices for a high-level understanding of the overall picture.

Results related to objective O3

In what follows, let the log-likelihood of the data stored in node k (using $D^{(k)}$) be denoted by

$$l^{(k)}(\beta, \phi) = \frac{1}{n^{(k)}} \sum_{i=1}^{n^{(k)}} \left\{ \frac{y_i^{(k)} h(\beta^\top x_i^{(k)}) - b\{h(\beta^\top x_i^{(k)})\}}{\phi} + c(y_i^{(k)}, \phi) \right\}.$$

Also, let $D^{(k)}(\beta) \in R^{p+1}$ be such that

$$D^{(k)}(\beta) = \frac{1}{n^{(k)}} \sum_{i=1}^{n^{(k)}} x_i^{(k)} h'(\beta^\top x_i^{(k)}) [y_i^{(k)} - b\{h(\beta^\top x_i^{(k)})\}], \quad (2)$$

and define the $(p+1) \times (p+1)$ matrix $V^{(k)}(\beta)$ as

$$V^{(k)}(\beta) = \frac{1}{n^{(k)}} \sum_{i=1}^{n^{(k)}} x_i^{(k)} (x_i^{(k)})^\top (h'(\beta^\top x_i^{(k)})^2 b''\{h(\beta^\top x_i^{(k)})\} - h''(\beta^\top x_i^{(k)}) [y_i^{(k)} - b\{h(\beta^\top x_i^{(k)})\}]). \quad (3)$$

Since $D^{(k)}(\beta) = \phi \nabla_\beta l^{(k)}(\beta, \phi)$, then, solving the equation $D^{(k)}(\beta) = 0$ yields the node-specific MLE of β , denoted hereafter by $\hat{\beta}_{MLE}^{(k)}$. The matrix $V^{(k)}(\beta)$ is equal to $-\nabla_\beta D^{(k)}(\beta)$ and relates to Fisher information matrix through the equation $V^{(k)}(\beta) = -\phi \nabla_\beta^2 l^{(k)}(\beta, \phi)$.

Finally, set

$$E^{(k)}(\phi, \beta) = \frac{1}{n^{(k)}} \sum_{i=1}^{n^{(k)}} \left[y_i^{(k)} h(\beta^\top x_i^{(k)}) - b\{h(\beta^\top x_i^{(k)})\} \right] - \frac{\phi^2}{n^{(k)}} \sum_{i=1}^{n^{(k)}} \frac{\partial}{\partial \phi} c(y_i^{(k)}, \phi), \quad (4)$$

and

$$F^{(k)}(\phi) = \frac{2\phi}{n^{(k)}} \sum_{i=1}^{n^{(k)}} \frac{\partial}{\partial \phi} c(y_i^{(k)}, \phi) + \frac{\phi^2}{n^{(k)}} \sum_{i=1}^{n^{(k)}} \frac{\partial^2}{\partial \phi^2} c(y_i^{(k)}, \phi). \quad (5)$$

Because $E^{(k)}(\phi, \beta) = -\phi^2 (\partial/\partial \phi) l^{(k)}(\beta, \phi)$, when ϕ is unknown, solving the equation $E^{(k)}(\phi, \beta_{MLE}^{(k)}) = 0$ for ϕ yields its node-specific MLE of ϕ^* . We have $F^{(k)}(\phi) = -\partial/\partial \phi E^{(k)}(\phi, \beta)$.

Simple averaging

The simple averaging method follows upon execution of Algorithm 1. First, each data node computes their local maximum by solving successively $D^{(k)}(\beta) = 0$ and $E^{(k)}(\phi, \hat{\beta}_{MLE}^{(k)}) = 0$. To compute the confidence intervals at the CC level, the entries of the $(p+1) \times (p+1)$ matrix $V_{MLE}^{(k)} = V^{(k)}(\hat{\beta}_{MLE}^{(k)})$ have to be computed from formula (3) with $\beta = \hat{\beta}_{MLE}^{(k)}$. Then, the set

$$S_0^{(k)} = [\hat{\beta}_{MLE}^{(k)}, \hat{\phi}_{MLE}^{(k)}, V_{MLE}^{(k)}] \quad (6)$$

is sent to the CC. The parameter estimates are then aggregated by the CC through averaging. Specifically, the CC computes

$$\hat{\beta}_{SA} = \sum_{k=1}^K w^{(k)} \hat{\beta}_{MLE}^{(k)} \text{ and } \hat{\phi}_{SA} = \sum_{k=1}^K w^{(k)} \hat{\phi}_{MLE}^{(k)}, \quad (7)$$

where $w^{(1)}, \dots, w^{(K)}$ are weights (i.e., $w^{(k)} \geq 0$ and $\sum_{k=1}^K w^{(k)} = 1$) used to combine each node's contribution.

Often, weights can be taken proportional to local sample sizes, leading to the choice $w^{(k)} = n^{(k)}/n$.

Wald-type confidence intervals for β^* can be constructed based on the fact that the sequence $\sqrt{n}(\hat{\beta}_{SA} - \beta^*)$ converges in distribution to a centred normal random variable with covariance matrix

$$\Sigma_{SA} = \phi^* \sum_{k=1}^K \frac{w^{(k)2}}{p^{(k)}} T_{\beta^*}^{(k)}, \text{ where } T_{\beta^*}^{(k)} = E\{V^{(k)}(\beta^*)\}.$$

See the Appendix B section Theory for the adapted simple averaging estimator. Since $T_{\beta^*}^{(k)}$ is consistently estimated by $V_{MLE}^{(k)}$ and ϕ^* by $\hat{\phi}_{SA}$, and as $p^{(k)}$ can be estimated by $n^{(k)}/n$, it follows that a consistent estimator for Σ_{SA} is given by

$$\hat{\Sigma}_{SA} = \hat{\phi}_{SA} \sum_{k=1}^K \frac{n w^{(k)2}}{n^{(k)}} V_{MLE}^{(k)}.$$

The simple averaging final estimates are then given by

$$R = [\hat{\beta}_{SA}, \hat{\Sigma}_{SA}].$$

Algorithm 1. Simple averaging inference procedure.

Input at the CC level: Weight $w^{(1)}, \dots, w^{(k)}$ **attributed to each node's contribution**

Step required from each node $k \in \{1, \dots, K\}$:

Using data in $D^{(k)}$, compute the following quantities:

- MLE $\hat{\beta}_{MLE}^{(k)}$ of β^* by solving $D^{(k)}(\beta)=0$;
- MLE $\hat{\phi}_{MLE}^{(k)}$ of ϕ^* by solving $E^{(k)}(\phi, \beta_{MLE}^{(k)})=0$;
- $V_{MLE}^{(k)} = V^{(k)}(\hat{\beta}_{MLE}^{(k)})$ using formula (3) with $\beta = \hat{\beta}_{MLE}^{(k)}$

Send to the CC: $S_0^{(k)} = \{\hat{\beta}_{MLE}^{(k)}, \hat{\phi}_{MLE}^{(k)}, V_{MLE}^{(k)}\}$.

Step required from the CC:

Using the received sets of quantities $S_0^{(1)}, \dots, S_0^{(K)}$, calculate

- the simple averaging estimators $\hat{\beta}_{SA} = \sum_{k=1}^K w^{(k)} \hat{\beta}_{MLE}^{(k)}$ and $\hat{\phi}_{SA} = \sum_{k=1}^K w^{(k)} \hat{\phi}_{MLE}^{(k)}$;
- the estimator of the variance-covariance matrix $\hat{\Sigma}_{SA} = \hat{\phi}_{SA} \sum_{k=1}^K \frac{n w^{(k)2}}{n^{(k)}} V_{MLE}^{(k)}$

Output from the CC:

Final estimates: $R = \{\hat{\beta}_{SA}, \hat{\Sigma}_{SA}\}$

Single distributed Newton-Raphson updating

The single distributed Newton-Raphson updating method follows upon execution of Algorithm 2 with $T=1$. First, the CC gathers summary statistics to compute the simple averaging estimators of β^* and ϕ^* without their accompanying confidence interval. Hence, for $k \in \{1, \dots, K\}$, and with $\hat{\beta}_{MLE}^{(k)}$ as above (6), node k sends to the CC the quantities

$$S_0^{(k)} = [\hat{\beta}_{MLE}^{(k)}, \hat{\phi}_{MLE}^{(k)}], \quad (8)$$

which uses them to compute $\hat{\beta}_{SA}$ and $\hat{\phi}_{SA}$ using the formulas in (7).

For reasons of convenience that will become clear later, the notation $\hat{\beta}_{NR,0}$ and $\hat{\phi}_{NR,0}$ will be utilized instead of $\hat{\beta}_{SA}$ and $\hat{\phi}_{SA}$, respectively. In this notation, the set of values

$$C_1 = \{\hat{\beta}_{NR,0}, \hat{\phi}_{NR,0}\}$$

is broadcasted to data nodes, which are then tasked with computing and sending back the quantities

$$S_1^{(k)} = [D_{NR,1}^{(k)}, V_{NR,1}^{(k)}, E_{NR,1}^{(k)}, F_{NR,1}^{(k)}],$$

where for any integer $t \geq 1$, one defines

$$\begin{aligned} D_{NR,t}^{(k)} &= D^{(k)}(\hat{\beta}_{NR,t-1}) \\ V_{NR,t}^{(k)} &= V^{(k)}(\hat{\beta}_{NR,t-1}) \\ E_{NR,t}^{(k)} &= E^{(k)}(\hat{\phi}_{NR,t-1}, \hat{\beta}_{NR,t-1}) \\ F_{NR,t}^{(k)} &= F^{(k)}(\hat{\phi}_{NR,t-1}, \hat{\beta}_{NR,t-1}). \end{aligned} \quad (9)$$

Upon receiving the $S_1^{(k)}$'s from each node, the CC calculates the following weighted averages:

$$\begin{aligned} \bar{D}_{NR,1} &= \sum_{k=1}^K w^{(k)} D_{NR,1}^{(k)}, \bar{V}_{NR,1} = \sum_{k=1}^K w^{(k)} V_{NR,1}^{(k)}, \\ \bar{E}_{NR,1} &= \sum_{k=1}^K w^{(k)} E_{NR,1}^{(k)}, \text{ and } \bar{F}_{NR,1} = \sum_{k=1}^K w^{(k)} F_{NR,1}^{(k)}. \end{aligned}$$

This enables the CC to execute Newton-Raphson updates from $\hat{\beta}_{NR,0}$ and $\hat{\phi}_{NR,0}$, respectively:

$$\begin{aligned} \hat{\beta}_{NR,1} &= \hat{\beta}_{NR,0} + \bar{V}_{NR,1}^{-1} \bar{D}_{NR,1} \\ \hat{\phi}_{NR,1} &= \hat{\phi}_{NR,0} + \bar{F}_{NR,1}^{-1} \bar{E}_{NR,1}. \end{aligned} \quad (10)$$

It is shown in the Appendix B section Theory for the adapted single distributed Newton-Raphson updating estimator that

$$\sqrt{n}(\hat{\beta}_{NR,1} - \beta^*) \rightarrow N(0, \Sigma_{NR})$$

Since $T_{\beta^*}^{(k)}$ is consistently estimated by $V_{NR,1}^{(k)}$ and ϕ^* by $\hat{\phi}_{NR,1}$, and as $p^{(k)}$ can be estimated by $n^{(k)}/n$, it follows that a consistent estimator for Σ_{NR} is given by

$$\hat{\Sigma}_{NR}$$

The method's final estimates are then given by

$$R = [\hat{\beta}_{NR,1}, \hat{\Sigma}_{NR}].$$

Algorithm 2. Distributed Newton-Raphson updatings algorithm.

Input at the CC level: Weight $w^{(1)}, \dots, w^{(K)}$ **attributed to each node's contribution**

Step required from each node $k \in \{1, \dots, K\}$:

Using data in $D^{(k)}$, compute

- $\hat{\beta}_{MLE}^{(k)}$ by solving $D^{(k)}(\beta) = 0$;
- $\hat{\phi}_{MLE}^{(k)}$ by solving $E^{(k)}(\phi, \hat{\beta}_{MLE}^{(k)}) = 0$

Send to the CC: $S_0^{(k)} = \{\hat{\beta}_{MLE}^{(k)}, \hat{\phi}_{MLE}^{(k)}\}$.

Step required from the CC:

Using the received quantities $S_0^{(1)}, \dots, S_0^{(K)}$:

- Calculate $\hat{\beta}_{SA} = \sum_{k=1}^K w^{(k)} \hat{\beta}_{MLE}^{(k)}$ and $\hat{\phi}_{SA} = \sum_{k=1}^K w^{(k)} \hat{\phi}_{MLE}^{(k)}$;
- Initialize $\hat{\beta}_{NR,t=0} = \hat{\beta}_{SA}$ and $\hat{\phi}_{NR,t=0} = \hat{\phi}_{SA}$.

Execute for $t=1, \dots, T$:

Step required from the CC:

Broadcast to nodes: $C_t = \{\hat{\beta}_{NR,t-1}, \hat{\phi}_{NR,t-1}\}$

Step required from each node $k \in \{1, \dots, K\}$:

Using data in $D^{(k)}$ and quantities in C_t , compute:

- $D_{NR,t}^{(k)}$ using formula (2) with $\beta = \hat{\beta}_{NR,t-1}$;
- $V_{NR,t}^{(k)}$ using formula (3) with $\beta = \hat{\beta}_{NR,t-1}$;
- $E_{NR,t}^{(k)}$ using formula (4) with $\phi = \hat{\phi}_{NR,t-1}$ and $\beta = \hat{\beta}_{NR,t-1}$;
- $F_{NR,t}^{(k)}$ using formula (5) with $\phi = \hat{\phi}_{NR,t-1}$ and $\beta = \hat{\beta}_{NR,t-1}$.

Send to the CC: $S_t^{(k)} = \{D_{NR,t}^{(k)}, V_{NR,t}^{(k)}, E_{NR,t}^{(k)}, F_{NR,t}^{(k)}\}$.

Step required from the CC:

Using the quantities in $S_t^{(k)}$, compute

- $\bar{D}_{NR,t} = \sum_{k=1}^K w^{(k)} D_{NR,t}^{(k)}$
- $\bar{V}_{NR,t} = \sum_{k=1}^K w^{(k)} V_{NR,t}^{(k)}$
- $\bar{E}_{NR,t} = \sum_{k=1}^K w^{(k)} E_{NR,t}^{(k)}$
- $\bar{F}_{NR,t} = \sum_{k=1}^K w^{(k)} F_{NR,t}^{(k)}$

Using, $\hat{\beta}_{NR,t-1}$, $\hat{\phi}_{NR,t-1}$ and the aggregated quantities, update previous parameter estimates

- $\hat{\beta}_{NR,t} = \hat{\beta}_{NR,t-1} + \bar{V}_{NR,t}^{-1} \bar{D}_{NR,t}$
- $\hat{\phi}_{NR,t} = \hat{\phi}_{NR,t-1} + \frac{\bar{E}_{NR,t}}{\bar{F}_{NR,t}}$

Step required from the CC:

Compute $\hat{\Sigma}_{NR} = \hat{I}$.

Output from the CC:

Estimates $R = \{\hat{\beta}_{NR,T}, \hat{\Sigma}_{NR}\}$

Multiple distributed Newton-Raphson updatings

The multiple distributed Newton-Raphson updatings method follows upon execution of Algorithm 2 with $T > 1$.

The first communication cycle follows the same procedure as described above for the single distributed Newton-Raphson updating method. It involves distributively computing a simple-averaging estimator and then performing a Newton-Raphson iteration starting from this estimator. The Newton descent is calculated as described in equation (10).

Formally, the algorithm begins with each data node k sending the set of quantities $S_0^{(k)}$ as described in equation (8) to the CC. Next, the CC calculates the simple averaging estimators using formula (7), and uses them to initialize $\hat{\beta}_{NR, \text{Step}=0} = \hat{\beta}_{SA}$ and $\hat{\phi}_{NR, \text{Step}=0} = \hat{\phi}_{SA}$.

The following steps are then repeated for a certain number of iterations. At iteration t , starting from $t=1$, the CC broadcasts the values $C_t = (\hat{\beta}_{NR, t-1}, \hat{\phi}_{NR, t-1})$ to the data nodes. The data nodes compute the quantities $D_{NR, t}^{(k)}$, $V_{NR, t}^{(k)}$, $E_{NR, t}^{(k)}$, and $F_{NR, t}^{(k)}$ as defined in equation (9), and send them back to the CC.

The CC then utilizes these quantities to perform a Newton update. Specifically, it calculates $\hat{\beta}_{NR, t} = \hat{\beta}_{NR, t-1} + V_{NR, t}^{-1} D_{NR, t}$ and $\hat{\phi}_{NR, t} = \hat{\phi}_{NR, t-1} + E_{NR, t} / F_{NR, t}$.

If the iterative cycle is repeated until convergence, the resulting estimates of β^* are equivalent to the maximum likelihood estimators derived from pooled data. This is because, in GLMs, for maximum likelihood estimators, if both the pooled and distributed algorithms are initialized with the same values for $\hat{\beta}_{NR, \text{Step}=0}$ and $\hat{\phi}_{NR, \text{Step}=0}$, then at each subsequent iteration, the distributed Newton update computed by the CC will be identical to the update obtained in a pooled setting.

For the method to yield consistent estimates, it is not necessary to initialize it with simple averaging estimators. However, using simple averaging estimators as initialization may speed up convergence, particularly in large sample sizes, since these estimators are \sqrt{n} -consistent.

Let $\hat{\beta}_{NR, T}$ denote the estimator obtained at convergence. Since it is (nearly) equal to the pooled MLE of β^* , we can deduce from the Appendix B section Theory for the pooled centralized setting estimator that

$$\sqrt{n}(\hat{\beta}_{NR, T} - \beta^*) \rightarrow N(0, \Sigma_{NR}) \quad (11)$$

Following the same reasoning used earlier for the single distributed Newton-Raphson update method, we can consistently estimate the variance-covariance matrix as

$$\hat{\Sigma}_{NR}$$

Distributed estimating equation

The distributed estimating equation algorithm follows upon execution of Algorithm 3. First, each node is responsible for computing the MLEs $\hat{\beta}_{MLE}^{(k)}$ and $\hat{\phi}_{MLE}^{(k)}$ of β^* and ϕ^* , respectively, using its own data. These estimators, along with the hessian matrix $V_{MLE}^{(k)} = V^{(k)}(\hat{\beta}_{MLE}^{(k)})$ and $F_{MLE}^{(k)} = F^{(k)}$, are then sent to the CC. The set

$$S_0^{(k)} = (\hat{\beta}_{MLE}^{(k)}, \hat{\phi}_{MLE}^{(k)}, V_{MLE}^{(k)}, F_{MLE}^{(k)}) \quad (12)$$

is transmitted to the CC. The CC calculates the weighted average of the Hessians and the $F_{MLE}^{(k)}$ values as follows:

$$\bar{V}_{EE} = \sum_{k=1}^K w^{(k)} V_{MLE}^{(k)} \text{ and } \bar{F}_{EE} = \sum_{k=1}^K w^{(k)} F_{MLE}^{(k)}. \quad (13)$$

The parameter estimates can then be calculated as

$$\hat{\beta}_{EE} = \bar{V}_{EE}^{-1} \sum_{k=1}^K w^{(k)} V_{MLE}^{(k)} \hat{\beta}_{MLE}^{(k)} \quad \hat{\phi}_{EE} = \bar{F}_{EE}^{-1} \sum_{k=1}^K w^{(k)} F_{MLE}^{(k)} \hat{\phi}_{MLE}^{(k)}. \quad (14)$$

It is shown in the Appendix B section Theory for the distributed estimating equations estimator that $\sqrt{n}(\hat{\beta}_{EE} - \beta^*)$ converges in distribution to a centred normal random variable with variance-covariance matrix given by

$$\Sigma_{EE} = \mathcal{I}$$

It can be consistently estimated by

$$\hat{\Sigma}_{EE} = \mathcal{I}$$

Algorithm 3. Distributed estimating equations inference procedure.

Input at the CC level: Weight $w^{(1)}, \dots, w^{(k)}$ attributed to each node's contribution

Step required from each node $k \in \{1, \dots, K\}$:

Using data in $D^{(k)}$, compute the following quantities:

- MLE $\hat{\beta}_{MLE}^{(k)}$ of β^* by solving $D^{(k)}(\beta) = 0$;
- MLE $\hat{\phi}_{MLE}^{(k)}$ of ϕ^* by solving $E^{(k)}(\phi, \beta_{MLE}^{(k)}) = 0$;
- $V_{MLE}^{(k)} = V^{(k)}(\hat{\beta}_{MLE}^{(k)})$ using formula (3) with $\beta = \hat{\beta}_{MLE}^{(k)}$;
- $F_{MLE}^{(k)} = F^{(k)}(\hat{\phi}_{MLE}^{(k)})$ using formula (5) with $\phi = \hat{\phi}_{MLE}^{(k)}$.

Send to the CC: $S_0^{(k)} = \{\hat{\beta}_{MLE}^{(k)}, \hat{\phi}_{MLE}^{(k)}, V_{MLE}^{(k)}, F_{MLE}^{(k)}\}$.

Step required from the CC:

Using the received sets of quantities $S_0^{(1)}, \dots, S_0^{(K)}$, calculate

- Aggregated quantities $\bar{V}_{EE} = \sum_{k=1}^K w^{(k)} \hat{V}_{MLE}^{(k)}$ and $\bar{F}_{EE} = \sum_{k=1}^K w^{(k)} F_{MLE}^{(k)}$
- EE estimators $\hat{\beta}_{EE} = \bar{V}_{EE}^{-1} \sum_{k=1}^K w^{(k)} V_{MLE}^{(k)} \hat{\beta}_{MLE}^{(k)}$ and $\hat{\phi}_{EE} = \bar{F}_{EE}^{-1} \sum_{k=1}^K w^{(k)} F_{MLE}^{(k)} \hat{\phi}_{MLE}^{(k)}$;

the variance-covariance matrix $\hat{\Sigma}_{EE} = \hat{I}$.

Output from the CC:

Parameter estimates and CIs $R = \{\hat{\beta}_{SA}, \hat{\Sigma}_{EE}\}$

Distributed estimation using a single gradient-enhanced log-likelihood

This method operates through Algorithm 4. First, the necessary information is collected by the CC to compute the initial estimates of β and ϕ , denoted as $\hat{\beta}_{\text{SGE},0}$ and $\hat{\phi}_{\text{SGE},0}$. In what follows, we assume these estimates are obtained using the simple averaging estimators calculated through Algorithm 1. Subsequently, the CC broadcasts $C_{1,1} = \{\hat{\beta}_{\text{SGE},0}, \hat{\phi}_{\text{SGE},0}\}$ to node $k \in \{2, \dots, K\}$. Each node is then requested to compute and transmit back the following quantities:

$$S_{1,1}^{(k)} = \{D_{\text{SGE},1}^{(k)}, V_{\text{SGE},1}^{(k)}, E_{\text{SGE},1}^{(k)}\},$$

with $D_{\text{SGE},1}^{(k)} = D^{(k)}(\hat{\beta}_{\text{SGE},0})$, $V_{\text{SGE},1}^{(k)} = V^{(k)}(\hat{\beta}_{\text{SGE},0})$ and $E_{\text{SGE},1}^{(k)} = E^{(k)}(\hat{\phi}_{\text{SGE},0}, \hat{\beta}_{\text{SGE},0})$.

The CC aggregates the $D^{(k)}$'s and the $E^{(k)}$'s using averaging by calculating

$$D_{\text{SGE},1} = \sum_{k=2}^K w^{(k)} D_{\text{SGE},1}^{(k)} \text{ and } E_{\text{SGE},1} = \sum_{k=2}^K w^{(k)} E_{\text{SGE},1}^{(k)}.$$

The $V^{(k)}$'s are momentarily stored and will be used later to compute the estimator for the asymptotic variance-covariance matrix of the final estimator of β^* . The quantities

$$C_{2,1} = \{D_{\text{SGE},1}, E_{\text{SGE},1}\}$$

are then sent to node $k=1$. Node $k=1$ computes the global average of the $D^{(k)}$'s by adding its own counterpart, i.e., it first computes

$$\bar{D}_{\text{SGE},1} = D_{\text{SGE},1} + w^{(1)} D_{\text{SGE},1}^{(1)}, \text{ and } \bar{E}_{\text{SGE},1} = E_{\text{SGE},1} + w^{(1)} E_{\text{SGE},1}^{(1)},$$

and then solves the surrogate likelihood function. Formally, it finds successively the value $\hat{\beta}_{\text{SGE},1}^{(1)}$ and $\hat{\phi}_{\text{SGE},1}^{(1)}$ that solves

$$\begin{aligned} D^{(1)}(\beta) + \bar{D}_{\text{SGE},1} - D_{\text{SGE},1}^{(1)} &= 0 \\ E^{(1)}(\phi, \hat{\beta}_{\text{SGE},1}^{(1)}) + \bar{E}_{\text{SGE},1} - E_{\text{SGE},1}^{(1)} &= 0 \end{aligned}$$

The results are sent back to the CC, along with $V_{\text{SGE},1}^{(1)}$, yielding

$$S_{2,1}^{(1)} = \{\hat{\beta}_{\text{SGE},1}^{(1)}, \hat{\phi}_{\text{SGE},1}^{(1)}, V_{\text{SGE},1}^{(1)}\}.$$

If simple averaging estimators for $\hat{\beta}_{\text{SGE},0}$ and $\hat{\phi}_{\text{SGE},0}$ are chosen, then $\sqrt{n}(\hat{\beta}_{\text{SGE},1} - \beta^*)$ converges in distribution to a mean-zero multivariate normal random variable with variance-covariance matrix given by

$$\Sigma_{\text{SGE},1} = \phi^* \left(T_{\beta^*}^{(1)} \right)^{-1} \left\{ \sum_{k=1}^K \frac{w^{(k)2}}{p^{(k)}} \left(A_{\beta^*}^{(k)} \right)^T T_{\beta^*}^{(k)} A_{\beta^*}^{(k)} \right\} \left(T_{\beta^*}^{(1)} \right)^{-1}$$

where $A_{\beta^*}^{(k)} = \sqrt{p^{(k)}} I_{p+1} + (T_{\beta^*}^{(1)} - T_{\beta^*}^{(k)}) (T_{\beta^*}^{(k)})^{-1}$,

with I_{p+1} the $p+1$ square identity matrix and $T_{\beta^*} = \sum_{k=1}^K w^{(k)} T_{\beta^*}^{(k)}$. See the Appendix B section Theory for the distributed estimation using a single gradient-enhanced log-likelihood. The latter can be consistently estimated by

$$\hat{\Sigma}_{\text{SGE},1} = \hat{\phi}_{\text{SGE},1} \left(V_{\text{SGE},1}^{(1)} \right)^{-1} \left\{ \sum_{k=1}^K \frac{w^{(k)2}}{n^{(k)}} n \left(\hat{A}_{\text{SGE},1}^{(k)} \right)^T V_{\text{SGE},1}^{(k)} \hat{A}_{\text{SGE},1}^{(k)} \right\} \left(V_{\text{SGE},1}^{(1)} \right)^{-1},$$

where

$$\hat{A}_{\text{SGE},1}^{(k)} = \sqrt{\frac{n^{(k)}}{n}} I_{p+1} + \{ V_{\text{SGE},1}^{(1)} - \hat{T}_{\beta^*} \} \left(V_{\text{SGE},1}^{(k)} \right)^{-1}$$

$$\text{with } \hat{T}_{\beta^*} = \sum_{k=1}^K w^{(k)} V_{\text{SGE},1}^{(k)}.$$

Remark 1. In the paper [25], where the method described above was originally proposed, the authors discuss a version in which the latter process is repeated multiple times. Their version

assumes that the data is uniformly and randomly split across nodes. Under this assumption, the resulting estimator of β^* is asymptotically equivalent to the pooled estimator, regardless of the number of iterations executed. This equivalence occurs because when the predictors' distribution is the same across nodes and the node sample sizes are equal, then $T_{\beta^*}^{(k)} \equiv T_{\beta^*}$ and $p^{(k)} \equiv 1/K$. By choosing $w^{(k)} = 1/K$, it follows that $A_{\beta^*}^{(k)} = I_{p+1}$, resulting in the following expression for $\Sigma_{\text{SGE},1}$: $\Sigma_{\text{SGE},1} = \phi^*(T_{\beta^*})^{-1}$. The variance-covariance matrix above is also the same as that of the simple averaging estimator in the setting of equal sample sizes and even predictor distributions. Consequently, at each iteration, the probability distribution of the resulting estimator remains unchanged. However, in a more general setting where predictor distributions and sample sizes vary across nodes, these cancellations no longer occur. Therefore, in this case, the probability distribution of the obtained estimator changes after each iteration, and tracking these changes falls beyond the scope of objective 3, see the Appendix B section Theory for the distributed estimation using a single gradient-enhanced log-likelihood. Hence, the current presentation focused on the case where only one iteration is executed.

Algorithm 4. Inference procedure based on the distributed estimation using a single gradient-enhanced log-likelihood method.

Input at the CC level: Weight $w^{(1)}, \dots, w^{(k)}$ **attributed to each node's contribution**

Step required from each node $k \in \{1, \dots, K\}$:

Using data in $D^{(k)}$, compute

- $\hat{\beta}_{MLE}^{(k)}$ by solving $D^{(k)}(\beta) = 0$;
- $\hat{\phi}_{MLE}^{(k)}$ by solving $E^{(k)}(\phi, \hat{\beta}_{MLE}^{(k)}) = 0$

Send to the CC: $S_0^{(k)} = \{\hat{\beta}_{MLE}^{(k)}, \hat{\phi}_{MLE}^{(k)}\}$.

Step required from the CC:

Using the received sets of quantities $S_0^{(1)}, \dots, S_0^{(K)}$, calculate

- simple averaging estimators $\hat{\beta}_{SA} = \sum_{k=1}^K w^{(k)} \hat{\beta}_{MLE}^{(k)}$ and $\hat{\phi}_{SA} = \sum_{k=1}^K w^{(k)} \hat{\phi}_{MLE}^{(k)}$
- Initialize $\hat{\beta}_{SGE,0} = \hat{\beta}_{SA}$ and $\hat{\phi}_{SGE,0} = \hat{\phi}_{SA}$.

Broadcast to nodes: $C_{1,1} = \{\hat{\beta}_{SGE,0}, \hat{\phi}_{SGE,0}\}$.

Step required from each node $k \in \{2, \dots, K\}$:

Using data in $D^{(k)}$, compute the following quantities:

- $D_{SGE,1}^{(k)} = D^{(k)}(\hat{\beta}_{SGE,0})$ using formula (2) with $\beta = \hat{\beta}_{SGE,0}$;
- $V_{SGE,1}^{(k)} = V^{(k)}(\hat{\beta}_{SGE,0})$ using formula (3) with $\beta = \hat{\beta}_{SGE,0}$;
- $E_{SGE,1}^{(k)} = E^{(k)}(\hat{\phi}_{SGE,0}, \hat{\beta}_{SGE,0})$ using formula (4) with $(\phi, \beta) = (\hat{\phi}_{SGE,0}, \hat{\beta}_{SGE,0})$.

Send to the CC: $S_{1,1}^{(k)} = \{D_{SGE,1}^{(k)}, V_{SGE,1}^{(k)}, E_{SGE,1}^{(k)}\}$

Step required from the CC:

Using the received sets of quantities $S_{1,1}^{(2)}, \dots, S_{1,1}^{(K)}$, calculate

- $D_{SGE,1} = \sum_{k=2}^K w^{(k)} D_{SGE,1}^{(k)}$
- $E_{SGE,1} = \sum_{k=2}^K w^{(k)} E_{SGE,1}^{(k)}$

Broadcast to node $k=1$: $C_{2,1} = \{D_{SGE,1}, E_{SGE,1}\}$.

Step required from node $k=1$:

Using data in $D^{(1)}$, calculate

- $D_{SGE,1}^{(1)} = D^{(1)}(\hat{\beta}_{SGE,0})$ using formula (2) with $\beta = \hat{\beta}_{SGE,0}$;
- $V_{SGE,1}^{(1)} = V^{(1)}(\hat{\beta}_{SGE,0})$ using formula (3) with $\beta = \hat{\beta}_{SGE,0}$;
- $E_{SGE,1}^{(1)} = E^{(1)}(\hat{\phi}_{SGE,0}, \hat{\beta}_{SGE,0})$ using formula (4) with $(\phi, \beta) = (\hat{\phi}_{SGE,0}, \hat{\beta}_{SGE,0})$.
- $\bar{D}_{SGE,1} = D_{SGE,1} + w^{(1)} D_{SGE,1}^{(1)}$
- $\bar{E}_{SGE,1} = E_{SGE,1} + w^{(1)} E_{SGE,1}^{(1)}$

Send to the CC: $S_{2,1}^{(1)} = \{\hat{\beta}_{SGE,1}, \hat{\phi}_{SGE,1}, V_{SGE,1}^{(1)}\}$

Step required from the CC:

Compute

- $\hat{A}_{SGE,1}^{(k)} = w^{(k)} I_{p+1} + \left(V_{SGE,1}^{(1)} - \left(\sum_{k=1}^K w^{(k)} V_{SGE,1}^{(k)} \right) \right) \left(V_{SGE,1}^{(k)} \right)^{-1}$ for $k \in \{1, \dots, K\}$

$$- \hat{\Sigma}_{\text{SGE},1} = \hat{\phi}_{\text{SGE},1} \left(V_{\text{SGE},1}^{(1)} \right)^{-1} \left\{ \sum_{k=1}^K \frac{n}{n^{(k)}} \left(\hat{A}_{\text{SGE},1}^{(k)} \right)^T V_{\text{SGE},1}^{(k)} \hat{A}_{\text{SGE},1}^{(k)} \right\} \left(V_{\text{SGE},1}^{(1)} \right)^{-1}$$

Output from the CC:

Parameter estimates $R = \{ \hat{\beta}_{\text{SGE},1}, \hat{\Sigma}_{\text{SGE},1} \}$

Distributed estimation using multiple gradient-enhanced log-likelihood

This method operates through Algorithm 5. First, the CC collects the necessary information to compute the initial estimates, denoted as $\hat{\beta}_{\text{MGE},0}$ and $\hat{\phi}_{\text{MGE},0}$. In this case, we assume that these estimates are obtained using the simple averaging estimators calculated through Algorithm 1.

Subsequently, the CC broadcasts $C_{1,1} = \{\hat{\beta}_{\text{MGE},0}, \hat{\phi}_{\text{MGE},0}\}$ to each node, which is then requested to compute and transmit back the following quantities:

$$S_{1,1}^{(k)} = \{D_{\text{MGE},1}^{(k)}, V_{\text{MGE},1}^{(k)}, E_{\text{MGE},1}^{(k)}\}.$$

Here, $D_{\text{MGE},1}^{(k)} = D^{(k)}(\hat{\beta}_{\text{MGE},0})$, $V_{\text{MGE},1}^{(k)} = V^{(k)}(\hat{\beta}_{\text{MGE},0})$ and $E_{\text{MGE},1}^{(k)} = E^{(k)}(\hat{\phi}_{\text{MGE},0}, \hat{\beta}_{\text{MGE},0})$.

The CC aggregates the $D^{(k)}$'s and the $E^{(k)}$'s using averaging by calculating

$$\bar{D}_{\text{MGE},1} = \sum_{k=1}^K w^{(k)} D_{\text{MGE},1}^{(k)} \text{ and } \bar{E}_{\text{MGE},1} = \sum_{k=1}^K w^{(k)} E_{\text{MGE},1}^{(k)}.$$

The CC then broadcasts $C_{2,1} = \{\bar{D}_{\text{MGE},1}, \bar{E}_{\text{MGE},1}\}$ to each node, which are then tasked with solving the surrogate likelihood function. Formally, they find successively the value $\hat{\beta}_{\text{MGE},1}^{(k)}$ and $\hat{\phi}_{\text{MGE},1}^{(k)}$ that solves

$$\begin{aligned} D^{(k)}(\beta) + \bar{D}_{\text{MGE},1} - D_{\text{MGE},1}^{(k)} &= 0 \\ E^{(k)}(\phi, \hat{\beta}_{\text{MGE},1}^{(1)}) + \bar{E}_{\text{MGE},1} - E_{\text{MGE},1}^{(k)} &= 0 \end{aligned}$$

Each node then transmits their set of local surrogate likelihood estimators to the CC:

$$S_{2,1}^{(k)} = \{\hat{\beta}_{\text{MGE},1}^{(k)}, \hat{\phi}_{\text{MGE},1}^{(k)}\}.$$

Using the received sets of quantities $S_{2,1}^{(1)}, \dots, S_{2,1}^{(K)}$, the CC aggregates them through averaging using the following formulas:

$$\hat{\beta}_{\text{MGE},1} = \sum_{k=1}^K w^{(k)} \hat{\beta}_{\text{MGE},1}^{(k)} \text{ and } \hat{\phi}_{\text{MGE},1} = \sum_{k=1}^K w^{(k)} \hat{\phi}_{\text{MGE},1}^{(k)}.$$

It is shown in the Appendix B section Theory for the distributed estimation using multiple gradient-enhanced log-likelihoods that $\sqrt{n}(\hat{\beta}_{\text{MGE},1} - \beta^*)$ converges in distribution to a multivariate normal random variable with mean 0 and a variance-covariance matrix given by:

$$\begin{aligned} \Sigma_{\text{MGE},1} &= \phi^* \sum_{k=1}^K \frac{w^{(k)2}}{p^{(k)}} \left[\sqrt{p^{(k)}} U_{\beta^*} T_{\beta^*}^{(k)} + (I_{p+1} - U_{\beta^*} T_{\beta^*}) \right] \\ &\quad \times \left[\sqrt{p^{(k)}} U_{\beta^*} + (I_{p+1} - U_{\beta^*} T_{\beta^*}) (T_{\beta^*}^{(k)})^{-1} \right], \end{aligned}$$

where $U_{\beta^*} = \sum_{k=1}^K w^{(k)} (T_{\beta^*}^{(k)})^{-1}$. The latter can be consistently estimated with

$$\hat{\Sigma}_{\text{MGE},1} = \hat{\phi}_{\text{MGE},1} \sum_{k=1}^K \frac{n w^{(k)2}}{n^{(k)}} \left[\sqrt{\frac{n^{(k)}}{n}} \hat{U}_{\beta^*} V_{\text{MGE},1}^{(k)} + (I_{p+1} - \hat{U}_{\beta^*} \hat{T}_{\beta^*}) \right]$$

where $\hat{T}_{\beta^*} = \sum_{k=1}^K w^{(k)} V_{\text{MGE},1}^{(k)}$ and $\hat{U}_{\beta^*} = \sum_{k=1}^K w^{(k)} (V_{\text{MGE},1}^{(k)})^{-1}$.

Algorithm 5. Inference procedure based on the distributed estimation using a single gradient-enhanced log-likelihood method.

Input at the CC level: Weight $w^{(1)}, \dots, w^{(k)}$ **attributed to each node's contribution**

Step required from each node $k \in \{1, \dots, K\}$:

Using data in $D^{(k)}$, compute

- $\hat{\beta}_{MLE}^{(k)}$ by solving $D^{(k)}(\beta) = 0$;
- $\hat{\phi}_{MLE}^{(k)}$ by solving $E^{(k)}(\phi, \hat{\beta}_{MLE}^{(k)}) = 0$

Send to the CC: $S_0^{(k)} = \{\hat{\beta}_{MLE}^{(k)}, \hat{\phi}_{MLE}^{(k)}\}$.

Step required from the CC:

Using the received sets of quantities $S_0^{(1)}, \dots, S_0^{(K)}$, calculate

- simple averaging estimators $\hat{\beta}_{SA} = \sum_{k=1}^K w^{(k)} \hat{\beta}_{MLE}^{(k)}$ and $\hat{\phi}_{SA} = \sum_{k=1}^K w^{(k)} \hat{\phi}_{MLE}^{(k)}$
- Initialize $\hat{\beta}_{MGE,0} = \hat{\beta}_{SA}$ and $\hat{\phi}_{MGE,0} = \hat{\phi}_{SA}$.

Broadcast to nodes: $C_{1,1} = \{\hat{\beta}_{MGE,0}, \hat{\phi}_{MGE,0}\}$.

Step required from each node $k \in \{1, \dots, K\}$:

Using data in $D^{(k)}$, calculate

- $D_{MGE,1}^{(k)} = D^{(k)}(\hat{\beta}_{MGE,0})$ using formula (2) with $\beta = \hat{\beta}_{MGE,0}$;
- $V_{MGE,1}^{(k)} = V^{(k)}(\hat{\beta}_{MGE,0})$ using formula (3) with $\beta = \hat{\beta}_{MGE,0}$;
- $E_{MGE,1}^{(k)} = E^{(k)}(\hat{\phi}_{MGE,0}, \hat{\beta}_{MGE,0})$ using formula (4) with $(\phi, \beta) = (\hat{\phi}_{MGE,0}, \hat{\beta}_{MGE,0})$.

Send to the CC: $S_{1,1}^{(k)} = \{D_{MGE,1}^{(k)}, V_{MGE,1}^{(k)}, E_{MGE,1}^{(k)}\}$

Step required from the CC:

Using the received sets of quantities $S_{1,1}^{(1)}, \dots, S_{1,1}^{(K)}$, calculate

- $\bar{D}_{MGE,1} = \sum_{k=1}^K w^{(k)} D_{MGE,1}^{(k)}$
- $\bar{E}_{MGE,1} = \sum_{k=1}^K w^{(k)} E_{MGE,1}^{(k)}$

Broadcast to nodes: $C_{2,1} = \{\bar{D}_{MGE,1}, \bar{E}_{MGE,1}\}$.

Step required from each node $k \in \{1, \dots, K\}$:

Using data in $D^{(k)}$, calculate

- $\hat{\beta}_{MGE,1}^{(k)}$ that solves $D^{(k)}(\beta) + \bar{D}_{MGE,1} - D_{MGE,1}^{(k)} = 0$
- $\hat{\phi}_{MGE,1}^{(k)}$ that solves $E^{(k)}(\phi, \hat{\beta}_{MGE,1}^{(k)}) + \bar{E}_{MGE,1} - E_{MGE,1}^{(k)} = 0$

Send to the CC: $S_{2,1}^{(k)} = \{\hat{\beta}_{MGE,1}^{(k)}, \hat{\phi}_{MGE,1}^{(k)}\}$

Step required from the CC:

Compute

- $\hat{\beta}_{MGE,1} = \sum_{k=1}^K w^{(k)} \hat{\beta}_{MGE,1}^{(k)}$

$$\begin{aligned}
- \quad & \hat{\phi}_{\text{MGE},1} = \sum_{k=1}^K w^{(k)} \hat{\phi}_{\text{MGE},1}^{(k)} \\
- \quad & \hat{T}_{\beta'} = \sum_{k=1}^K w^{(k)} V_{\text{MGE},1}^{(k)} \\
- \quad & \hat{U}_{\beta'} = \sum_{k=1}^K w^{(k)} \left(V_{\text{MGE},1}^{(k)} \right)^{-1} \\
- \quad & \hat{\Sigma}_{\text{MGE},1} = \hat{\phi}_{\text{MGE},1} \sum_{k=1}^K \frac{n w^{(k)2}}{n^{(k)}} \hat{\mathbf{I}} \times \hat{\mathbf{I}}
\end{aligned}$$

Output from the CC:

Parameter estimates $R = \{ \hat{\beta}_{\text{MGE},1}, \hat{\Sigma}_{\text{MGE},1} \}$

Summary of quantities exchanged for the adapted methods

The following table presents a summary of the quantities exchanged between the nodes and the CC in both directions. Table 2 demonstrates that the quantities involved in exchanges from the nodes to the CC consist of parameter estimates, gradients ($D^{[k]}$ vectors), Hessians ($V^{[k]}$ matrices), as well as real numbers ($E^{[k]}$ and $F^{[k]}$). On the other hand, the quantities shared from the CC to the nodes primarily consist of parameter estimates. Notably, Methods 5 and 6 differentiate themselves by requiring the sharing of aggregated gradient vectors and Hessian matrices as well.



Table 2. Quantities shared in each adapted method's communication workflow.

Method	Exchanged quantities from nodes to CC		Exchanged quantities from CC to nodes
	$S_0^{[k]}$	$S_t^{[k]}, t \geq 1$	C_t
1. Simple averaging	$\hat{\beta}_{MLE}^{[k]} ; \hat{\phi}_{MLE}^{[k]} ; V_{MLE}^{[k]}$	-	-
2. Single distributed Newton-Raphson updating	$\hat{\beta}_{MLE}^{[k]} ; \hat{\phi}_{MLE}^{[k]}$	$D_{NR,1}^{[k]} ; V_{NR,1}^{[k]} ; E_{NR,1}^{[k]} ; F_{NR,1}^{[k]}$	$\hat{\beta}_{NR,0} ; \hat{\phi}_{NR,0}$
3. Multiple distributed Newton-Raphson updatings (with T Newton-Raphson updatings)	$\hat{\beta}_{MLE}^{[k]} ; \hat{\phi}_{MLE}^{[k]}$	$D_{NR,t}^{[k]} ; V_{NR,t}^{[k]} ; E_{NR,t}^{[k]} ; F_{NR,t}^{[k]}$	$\hat{\beta}_{NR,t-1} ; \hat{\phi}_{NR,t-1}$
4. Distributed estimating equations	$\hat{\beta}_{MLE}^{[k]} ; \hat{\phi}_{MLE}^{[k]} ; V_{MLE}^{[k]} ; F_{MLE}^{[k]}$	-	-
5. Distributed single gradient-enhanced log-likelihood	$\hat{\beta}_{MLE}^{[k]} ; \hat{\phi}_{MLE}^{[k]}$	Nodes 2 to K: $D_{SGE,1}^{[k]} ; V_{SGE,1}^{[k]} ; E_{SGE,1}^{[k]}$ Node 1: $\hat{\beta}_{SGE,1} ; \hat{\phi}_{SGE,1} ; V_{SGE,1}^{[1]}$	$\hat{\beta}_{SGE,0} ; \hat{\phi}_{SGE,0}$ To node 1 only: $D_{SGE,1} ; E_{SGE,1}$
6. Distributed multiple gradient-enhanced log-likelihood	$\hat{\beta}_{MLE}^{[k]} ; \hat{\phi}_{MLE}^{[k]}$	$D_{MGE,1}^{[k]} ; V_{MGE,1}^{[k]} ; E_{MGE,1}^{[k]} ; \hat{\beta}_{MGE,1}^{[k]} ; \hat{\phi}_{MGE,1}^{[k]}$	$\hat{\beta}_{MGE,0} ; \hat{\phi}_{MGE,0} ; \bar{D}_{MGE,1} ; \bar{E}_{MGE,1}$

Comparison of adapted methods

Table 3 compares the main adapted HPSA methods on the quantities shared between the CC and nodes and the operational complexity of the procedures.

Methods 1 and 4 require only one communication from the data nodes to the CC and no communication back from the CC to the nodes. These so-called one-shot methods have the lowest operational complexity. Method 4 requires the additional quantity $F_{MLE}^{(k)}$ to be transmitted from each node to the CC.

Methods 2 and 3 perform Newton-Raphson updating using some initial estimator as a basis, usually the simple averaging estimator. While Method 2 requires this initial estimator to be \sqrt{n} -consistent, if T is large enough, any initial value will work for Method 3 (although convergence may be slower). Both methods require $D^{(k)}, V^{(k)}, E^{(k)}$ and $F^{(k)}$ to be evaluated and sent to the CC T times, with $T=1$ for Method 2. Compared to Method 1, Method 2 requires the additional quantities $D^{(k)}, E^{(k)}$ and $F^{(k)}$, and Method 3 further requires these quantities to be evaluated and communicated multiple times.

Method 5 relies on an approximation of the log-likelihood function. It requires an initial estimator, usually the simple averaging estimator. This approach treats node 1 differently, making it solve the surrogate log-likelihood using aggregates from the other nodes and its own data. The CC sends the initial estimator to each node, then requires them to evaluate $D^{(k)}, V^{(k)}$ and $E^{(k)}$ and send the result back to the CC once. It averages the results and then communicates them to node 1, which solves the surrogate log-likelihood and sends its results back to the CC.

Method 6 applies Method 5 to every node, making each node solve the surrogate log-likelihood function with its own data before averaging the resulting local estimators.

Table 3. Comparison of adapted methods.

Method	Information shared		Number of communications		Workflow
	From nodes to CC	From CC to nodes	From nodes to CC	From CC to nodes	
1. Simple averaging	Local parameter estimates; Hessian matrix of log-likelihood (with respect to β only)	None	1	0	I in Figure 2
2. Single distributed Newton-Raphson updating	Local parameter estimates; Gradient and Hessian of log-likelihood	Simple averaging aggregated estimates of parameters	2	1	II in Figure 3 with $T=1$
3. Multiple distributed Newton-Raphson updatings (with T Newton-Raphson updatings)	Local parameter estimates; $T \times$ Gradient and Hessian of log-likelihood	Simple averaging aggregated estimates of parameters; $(T-1) \times$ Newton-updated parameter estimates	$T+1$	T	II in Figure 3 with $T>1$
4. Distributed estimating equations	Local parameter estimates; Hessian of log-likelihood	None	1	0	I in Figure 2
5. Distributed single gradient-enhanced log-likelihood	From all nodes: Local parameter estimates; Hessian of log-likelihood (with respect to β only) From nodes 2 to K: Gradient of log-likelihood; From node 1	To all nodes: Simple averaging aggregated estimates of parameters; To node 1 only: Average of local gradients and Hessians	All nodes: 2	Nodes 2 to K: 1 Node 1: 2	III in Figure 4 with $T=1$

		only: Gradient-enhanced parameter estimates				
6. Distributed multiple gradient- enhanced log- likelihood		Local parameter estimates; Gradient of log-likelihood; Hessian of log-likelihood (with respect to β only); Gradient- enhanced parameter estimates	Simple averaging aggregated estimates of parameters; Average of local gradients and Hessians	3	2	IV in Figure 5 with $T=1$

Discussion

Summary of Findings

The first objective (O1) of this study aimed to identify and map the methodological approaches used and developed in the literature regarding HPSA. To achieve this, we conducted a scoping review, which included 41 articles following our protocol. These articles were categorized based on the types of models and communication schemes involved, as presented in Table 1. The analysis revealed that the majority of methods included in the scoping review focused on methodological settings associated with massive data. The communication schemes of these methods were demonstrated through Workflows I, II, III and IV.

The second objective (O2) of this study aimed to describe the approaches that can be employed for basic GLM regression analyses and identify the distributional assumptions they require. To accomplish this, we identified six approaches and classified them within Workflows I-IV, enabling their comparison in terms of operational communication protocols within a unified framework. However, a limitation of these methods is that they assume identical node sample sizes and node covariate distributions. This assumption reduces their suitability in settings commonly encountered in health analytics, where data collecting nodes are prone to generating different covariate distributions.

The third objective (O3) of this study was to present methods that relaxed these assumptions by adapting the approaches identified in O2 to the unequal sample sizes and non-identical covariate sample distribution setting. Additionally, we compared these methods in terms of the information shared and operational complexity. This involved adapting the quantities and estimators described in the original articles and deriving new asymptotic results with relaxed assumptions. We defined a unified framework to describe inference procedures utilizing these methods. This unified framework encompasses common hypotheses and notation and allows for both estimation and the construction of confidence intervals, providing detailed steps for both the data nodes and the CC.

Challenges and Opportunities

Work pertaining to O1 illustrated why it is so challenging for researchers and data custodians alike to find information regarding HPSA. While the HPSA literature is very recent (all included articles were published in 2010 or later), the literature is non-homogeneous, and it has not come to a consensus on nomenclature. No universal terminology exists, and different terms are used in the different fields developing and applying HPSA methods. Many specific methods introduced in applied contexts are special cases of more general methods which may or may not be cited. These characteristics make finding useful and efficient keywords arduous. This required adapting our search strategy.

This difficulty is compounded by the fact that statistical inference is not the main focus of most of the HPSA literature. The majority of published work is in the prediction, learning and optimization contexts. As a result, method assumptions are rarely discussed. This can be a problem when adapting these methods for inference. Furthermore, the methodological setting is often assumed to be in the massive data context where data is randomly distributed between nodes. This allows the authors to make strong assumptions on node sample sizes and covariate distributions which may be unrealistic in the confidential data complex where different data sources. These methods cannot be used directly for inference using confidential health data. While some work remains to be done when the structure of association between the covariate and the outcome is heterogeneous between nodes, we adapted widely used methods for when the distribution of covariate and sample sizes between nodes are not identical.

Table 1 illustrates how the majority of HPSA methods are focused on parametric models. Some work has also been done for semi-parametric and non-parametric regression, and some methods are

introduced outside of the regression framework (although they can also be applied to regression). Many methods do not require communication of quantities from the CC to the data nodes: they only require one transmission from the nodes to the CC. Given the lack of awareness around HPSA, starting by implementing lower operational complexity methods while providing useful results offers a promising path.

LHS aims to create a responsive and adaptive health environment where continuous learning loops involving evidence, practice, and data are leveraged, ideally in real-time, to generate knowledge during practice and use that knowledge to improve practice [52]. This requires analysing data from a variety of sources. HPSA methods offer a way to tackle the challenges that come with multi-jurisdictional data.

The methods can be implemented "manually" (e.g. via email exchanges), but platforms enabling semi-automated distributed fittings of statistical models have been proposed in the literature (e.g. [53]). As LHS work by continuously monitoring and analysing data, rather than by cross-sectional studies only, automated platforms are necessary. Ideally, these platforms should be able to link and standardize data sources from multiple jurisdictions, as well as perform HPSA. On the other hand, explicit descriptions of these methods' algorithms and the quantities exchanged are not always easily accessible and this complicates the evaluation of the tools by data custodians and researchers.

This is especially important since it is essential to clarify here that operating an HPSA algorithm does NOT ensure confidentiality in and of itself.

For example, it is known that sharing sample moments can compromise confidentiality. It can be shown that a set of n observations is uniquely determined by its first n sample moments [54]. This could prove problematic for methods that rely on sharing the first few moments of each node's sample, especially if number of observations is low, as the sample could be partially reconstructed by the CC.

The results presented here contribute to this objective by clarifying the workflows and quantities exchanged by each method. Nevertheless, further analysis of the confidentiality preserved by HPSA methods is needed to fully understand the risk associated with the sharing of summary statistics, especially as more rounds of communication between the CC and data nodes are completed. The framework of differential privacy (DP) has been used to guarantee the preservation of confidentiality in a few HPSA methods, but a wider application of DP to existing and popular methods has yet to be explored.

Funding

Health Data Research Network Canada, Natural Sciences and Engineering Research Council of Canada, Fonds de recherche du Québec - Nature et Technologie, the Chaire en informatique de la santé de l'Université de Sherbrooke and the Chaire MEIE Québec - Le numérique au service des systèmes de santé apprenants.

Acknowledgments

We would like to thank the GRIIS members who enriched this work via multiple conversations over the last few months and kept us going. We would also like to thank Pr. Kim McGrail for her very insightful comments on this work.

Conflicts of interest

The authors declare no conflict of interest.

Abbreviations

CC: Coordinating centre

GLM: Generalized linear model

HPSA: Horizontally Partitioned Statistical Analytics

ICES: Institute for Clinical Evaluative Sciences

LHS: Learning health systems

MCHP: Manitoba Centre for Health Policy

MLE: Maximum likelihood estimator

PACS: Picture archiving and communication system

References

1. Sinha BK, Hartung J, Knapp G. Statistical meta-analysis with applications. John Wiley & Sons; 2011. ISBN:9780470290897
2. Duan R, Boland MR, Liu Z, Liu Y, Chang HH, Xu H, Chu H, Schmid CH, Forrest CB, Holmes JH, Schuemie MJ, Berlin JA, Moore JH, Chen Y. Learning from electronic health records across multiple sites: A communication-efficient and privacy-preserving distributed algorithm. *Journal of the American Medical Informatics Association* 2020;27(3):376–385. doi:10.1093/jamia/ocz199
3. Gao Y, Liu W, Wang H, Wang X, Yan Y, Zhang R. A review of distributed statistical inference. *Statistical Theory and Related Fields* 2022 May;6(2):89–99. doi:10.1080/24754269.2021.1974158
4. Huo X, Cao S. Aggregated inference. *Wiley Interdisciplinary Reviews: Computational Statistics Wiley Online Library*; 2019;11(1):e1451. doi:10.1002/wics.1451
5. Agresti A. Foundations of linear and generalized linear models. John Wiley & Sons; 2015. ISBN:9781118730034
6. Arksey H, O'Malley L. Scoping studies: Towards a methodological framework. *International journal of social research methodology* 2005;8(1):19–32. doi:10.1080/1364557032000119616
7. Peters MD, Godfrey CM, Khalil H, McInerney P, Parker D, Soares CB. Guidance for conducting systematic scoping reviews. *JBIC Evidence Implementation LWW*; 2015;13(3):141–146. doi:10.1097/XEB.0000000000000050
8. Wohlin C. Guidelines for snowballing in systematic literature studies and a replication in software engineering. *Proceedings of the 18th international conference on evaluation and assessment in software engineering* 2014. p. 1–10. doi:10.1145/2601248.2601268
9. Jalali S, Wohlin C. Systematic literature studies: Database searches vs. Backward snowballing. *Proceedings of the ACM-IEEE international symposium on empirical software engineering and measurement* 2012. p. 29–38. doi:10.1145/2372251.2372257
10. Levac D, Colquhoun H, O'Brien KK. Scoping studies: Advancing the methodology. *Implementation science* 2010;5:1–9. doi:10.1186/1748-5908-5-69
11. Duan R, Luo C, Schuemie MJ, Tong J, Liang CJ, Chang HH, Boland MR, Bian J, Xu H, Holmes JH, Forrest CB, Morton SC, Berlin JA, Moore JH, Mahoney KB, Chen Y. Learning from local to global: An efficient distributed algorithm for modeling time-to-event data. *Journal of the American Medical Informatics Association* 2020;27(7):1028–1036. doi:10.1093/jamia/ocaa044
12. Rosenblatt J, Nadler B. On the optimality of averaging in distributed statistical learning. *Information and Inference* 2016;5(4):379–404. doi:10.1093/imaiai/iaw013
13. Banerjee M, Durot C, Sen B. Divide and conquer in nonstandard problems and the super-efficiency phenomenon. *The Annals of Statistics* 2019;47(2). doi:10.1214/17-AOS1633
14. Chen X, Xie M. A split-and-conquer approach for analysis of extraordinarily large data. *Statistica Sinica* 2014;24:1655–1684. doi:10.5705/ss.2013.088
15. Basiri S, Ollila E, Koivunen V. Robust, scalable, and fast bootstrap method for analyzing large scale data. *IEEE Transactions on Signal Processing* 2016;64(4):1007–1017. doi:10.1109/TSP.2015.2498121
16. Battey H, Fan J, Liu H, Lu J, Zhu Z. Distributed testing and estimation under sparse high dimensional models. *The Annals of Statistics* 2018;46(3):1352–1382. doi:10.1214/17-AOS1587
17. Fan J, Guo Y, Wang K. Communication-efficient accurate statistical estimation. *Journal of the American Statistical Association* 2021; doi:10.1080/01621459.2021.1969238
18. Guo G, Sun Y, Jiang X. A partitioned quasi-likelihood for distributed statistical inference.

- Computational Statistics 2020;35(4):1577–1596. doi:10.1007/s00180-020-00974-4
19. Lin N, Xi R. Aggregated estimating equation estimation. *Statistics and Its Interface* 2011;4(1):73–83. doi:10.4310/SII.2011.v4.n1.a8
 20. Zhang Y, Duchi JC, Wainwright MJ. Communication-efficient algorithms for statistical optimization. 2012 IEEE 51st IEEE Conference on Decision and Control (CDC) Maui, HI, USA: IEEE; 2012. p. 6792–6792. doi:10.1109/CDC.2012.6426691
 21. Chang C, Bu Z, Long Q. CEDAR: Communication efficient distributed analysis for regressions. *Biometrics* 2022; doi:10.1111/biom.13786
 22. Wu Y, Jiang X, Kim J, Ohno-Machado L. Grid binary LOGistic REGression (GLORE): Building shared models without sharing data. *Journal of the American Medical Informatics Association* 2012;19(5):758–764. doi:10.1136/amiajnl-2012-000862
 23. Hector EC, Song PX-K. Joint integrative analysis of multiple data sources with correlated vector outcomes. *Annals of Applied Statistics* 2022;16(3):1700–1717. doi:10.1214/21-AOAS1563
 24. Huang C, Huo X. A distributed one-step estimator. *Mathematical Programming* 2019;174(1):41–76. doi:10.1007/s10107-019-01369-0
 25. Jordan MI, Lee JD, Yang Y. Communication-efficient distributed statistical inference. *Journal of the American Statistical Association* 2019;114(526):668–681. doi:10.1080/01621459.2018.1429274
 26. Mozafari-Majd E, Koivunen V. Two-stage robust and sparse distributed statistical inference for large-scale data. *IEEE Transactions on Signal Processing* 2022;70:5351–5365. doi:10.1109/TSP.2022.3216704
 27. Mozafari-Majd E, Koivunen V. Robust variable selection and distributed inference using t-based estimators for large-scale data. 2020 28th european signal processing conference (EUSIPCO) 2021. p. 2453–2457. doi:10.23919/Eusipco47968.2020.9287773
 28. Yue X, Kontar RA, Gómez AME. Federated data analytics: A study on linear models. *IIE Transactions* 2022; doi:10.1080/24725854.2022.2157912
 29. Duan R, Ning Y, Chen Y. Heterogeneity-aware and communication-efficient distributed statistical inference. *Biometrika* 2022;109(1):67–83. doi:10.1093/biomet/asab007
 30. Tong J, Duan R, Li R, Scheuemie MJ, Moore JH, Chen Y. Robust-ODAL: Learning from heterogeneous health systems without sharing patient-level data. *Biocomputing 2020 Kohala Coast, Hawaii, USA: WORLD SCIENTIFIC; 2019. p. 695–706. doi:10.1142/9789811215636_0061*
 31. Di F, Wang L, Lian H. Communication-efficient estimation and inference for high-dimensional quantile regression based on smoothed decorrelated score. *Statistics in medicine* 2022;41(25):5084–5101. doi:10.1002/sim.9555
 32. Edmondson MJ, Luo C, Nazmul Islam M, Sheils NE, Buresh J, Chen Z, Bian J, Chen Y. Distributed quasi-poisson regression algorithm for modeling multi-site count outcomes in distributed data networks. *Journal of Biomedical Informatics* 2022;131. doi:10.1016/j.jbi.2022.104097
 33. Luo L, Li L. Online two-way estimation and inference via linear mixed-effects models. *Statistics in medicine* 2022;41(25):5113–5133. doi:10.1002/sim.9557
 34. Shu D, Young JG, Toh S. Privacy-protecting estimation of adjusted risk ratios using modified poisson regression in multi-center studies. *BMC Medical Research Methodology* 2019;19(1). doi:10.1186/s12874-019-0878-6
 35. Zhao T, Cheng G, Liu H. A partially linear framework for massive heterogeneous data. *The Annals of Statistics* 2016;44(4). doi:10.1214/15-AOS1410
 36. Park JA, Kim TH, Kim J, Park YR. WICOX: Weight-based integrated cox model for time-to-event data in distributed databases without data-sharing. *IEEE Journal of Biomedical and Health Informatics* 2023;27(1):526–537. doi:10.1109/JBHI.2022.3218585

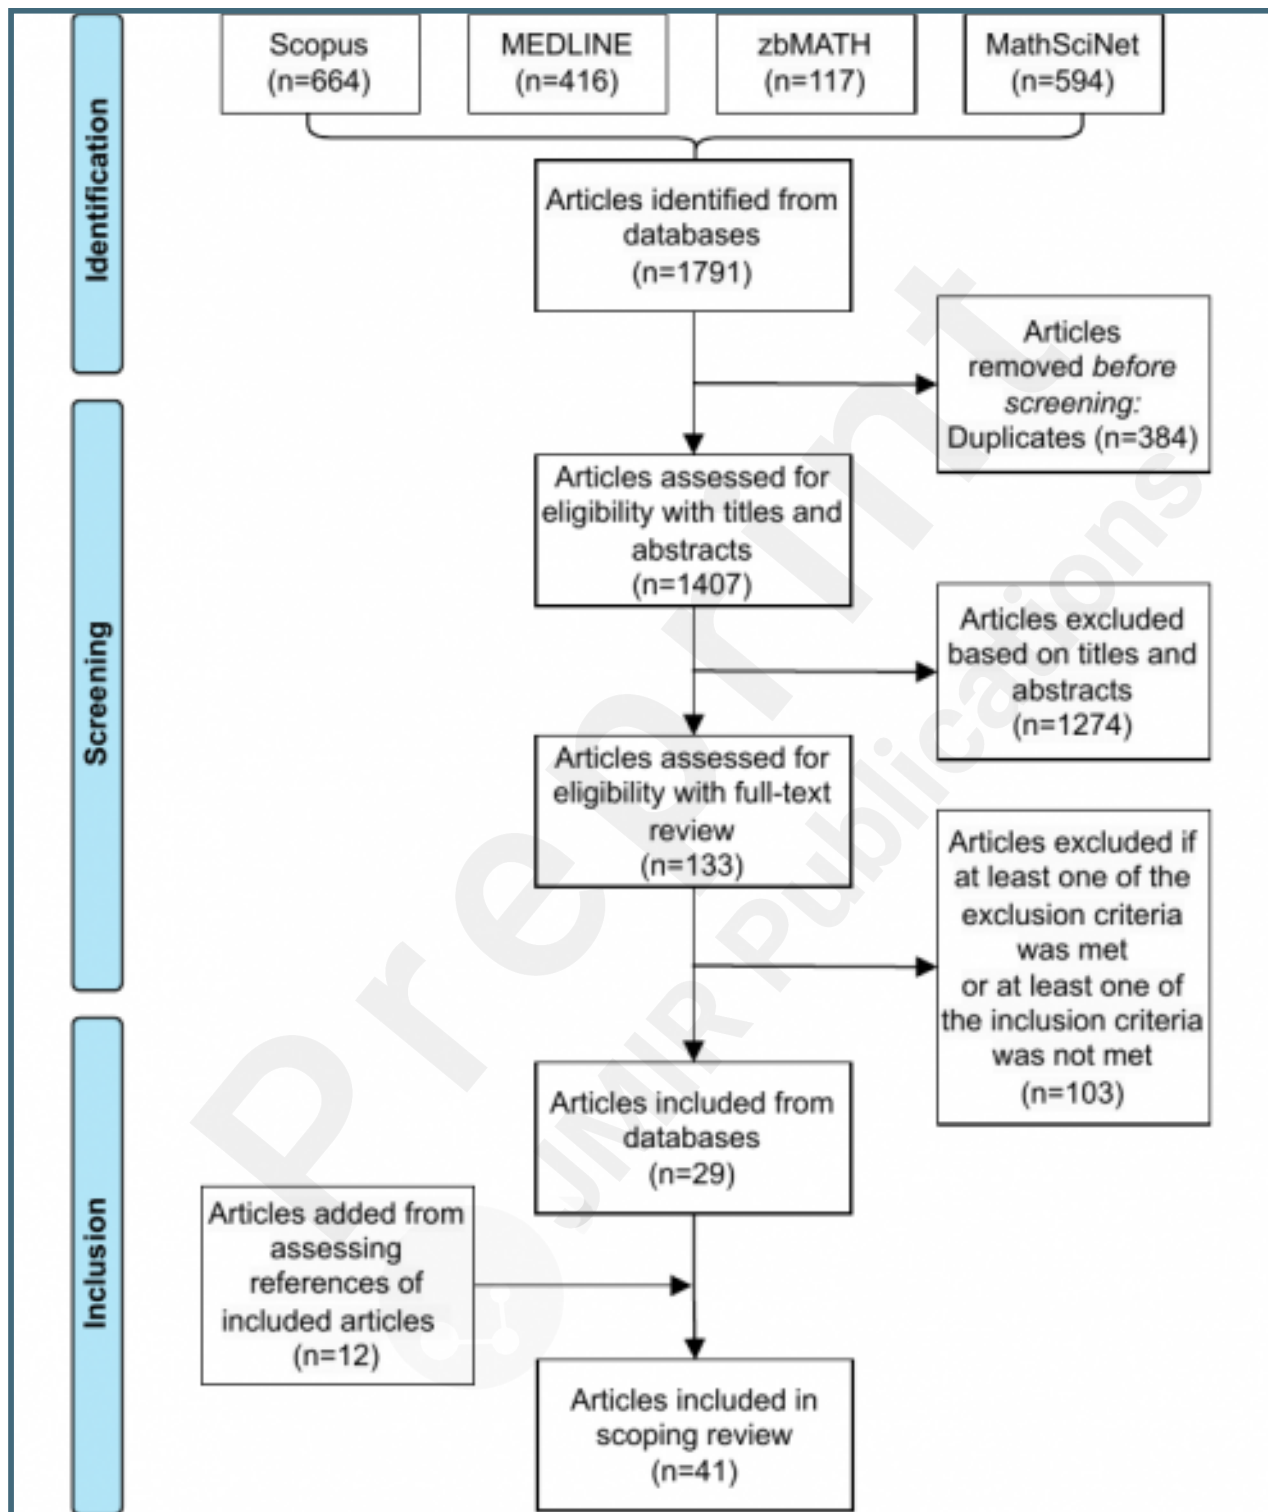
37. Luo J, Sun Q, Zhou W-X. Distributed adaptive huber regression. *Computational Statistics and Data Analysis* 2022;169. doi:10.1016/j.csda.2021.107419
38. Liu M, Shang Z, Cheng G. Nonparametric distributed learning under general designs. *Electronic Journal of Statistics* 2020;14(2):3070–3102. doi:10.1214/20-EJS1733
39. Zhang L, Castillo E del, Berglund AJ, Tingley MP, Govind N. Computing confidence intervals from massive data via penalized quantile smoothing splines. *Computational Statistics & Data Analysis* 2020;144:106885, –25. doi:10.1016/j.csda.2019.106885
40. Volgushev S, Chao S-K, Cheng G. Distributed inference for quantile regression processes. *Annals of Statistics* 2019;47(3):1634–1662. doi:10.1214/18-AOS1730
41. Wang X, Yang Z, Chen X, Liu W. Distributed inference for linear support vector machine. *Journal of Machine Learning Research* 2019;20. doi:10.48550/arXiv.1811.11922
42. Atta-Asiamah E, Yuan M. Distributed inference for degenerate u-statistics. *Stat* 2019;8(1). doi:10.1002/sta4.234
43. Minsker S. Distributed statistical estimation and rates of convergence in normal approximation. *Electronic Journal of Statistics* 2019;13(2):5213–5252. doi:10.1214/19-EJS1647
44. Lin N, Xi R. Fast surrogates of u-statistics. *Computational Statistics & Data Analysis* 2010;54(1):16–24. doi:10.1016/j.csda.2009.08.009
45. Bruce S, Li Z, Yang H-C, Mukhopadhyay S. Nonparametric distributed learning architecture for big data: Algorithm and applications. *IEEE Transactions on Big Data* 2019;5(2):166–179. doi:10.1109/TBDATA.2018.2810187
46. Chen SX, Peng L. Distributed statistical inference for massive data. *Annals of Statistics* 2021;49(5):2851–2869. doi:10.1214/21-AOS2062
47. Nezakati E, Pircalabelu E. Unbalanced distributed estimation and inference for the precision matrix in gaussian graphical models. *Statistics and Computing* 2023;33(2). doi:10.1007/s11222-023-10211-9
48. Shi C, Lu W, Song R. A massive data framework for m-estimators with cubic-rate. *Journal of the American Statistical Association* 2018;113(524):1698–1709. doi:10.1080/01621459.2017.1360779
49. Wu S, Xu Y, Feng Z, Yang X, Wang X, Gao X. Multiple-platform data integration method with application to combined analysis of microarray and proteomic data. *BMC Bioinformatics* 2012;13(1). doi:10.1186/1471-2105-13-320
50. Lai RCS, Hannig J, Lee TCM. Method g: Uncertainty quantification for distributed data problems using generalized fiducial inference. *Journal of Computational and Graphical Statistics* 2021;30(4):934–945. doi:10.1080/10618600.2021.1923514
51. Shamir O, Srebro N, Zhang T. Communication-efficient distributed optimization using an approximate newton-type method. *International conference on machine learning PMLR*; 2014. p. 1000–1008. doi:10.48550/arXiv.1312.7853
52. Guise JM, Savitz LA, Friedman CP. Mind the Gap: Putting Evidence into Practice in the Era of Learning Health Systems. *Journal of General Internal Medicine* 2018;33(12):2237–2239. doi: 10.1007/s11606-018-4633-1.
53. Beyan O, Choudhury A, Soest J van, Kohlbacher O, Zimmermann L, Stenzhorn H, Karim MdR, Dumontier M, Decker S, Silva Santos LOB da, Dekker A. Distributed Analytics on Sensitive Medical Data: The Personal Health Train. *Data Intelligence* 2020 Jan;2(1-2):96–107. doi:10.1162/dint_a_00032
54. Provost SB, Zareamoghaddam H, Ahmed SE, Ha H-T. The generalized pearson family of distributions and explicit representation of the associated density functions. *Communications in Statistics - Theory and Methods Taylor & Francis*; 2022;51(16):5590–5606. doi:10.1080/03610926.2020.1843680
55. Lu C-L, Wang S, Ji Z, Wu Y, Xiong L, Jiang X, Ohno-Machado L. WebDISCO: A web

- service for distributed cox model learning without patient-level data sharing. *Journal of the American Medical Informatics Association Oxford University Press*; 2015;22(6):1212–1219. doi:10.1093/jamia/ocv083
56. Toh S, Wellman R, Coley RY, Horgan C, Sturtevant J, Moyneur E, Janning C, Pardee R, Coleman KJ, Arterburn D, McTigue K, Anau J, Cook AJ. Combining distributed regression and propensity scores: A doubly privacy-protecting analytic method for multicenter research. *Clinical Epidemiology* 2018; Volume 10:1773–1786. doi:10.2147/CLEP.S178163
57. Xiong R, Koenecke A, Powell M, Shen Z, Vogelstein JT, Athey S. Federated Causal Inference in Heterogeneous Observational Data. *arXiv*; 2022. doi:10.48550/arXiv.2107.11732
58. Vo TV, Hoang TN, Lee Y, Leong T-Y. Federated Estimation of Causal Effects from Observational Data. *arXiv*; 2021. doi:10.48550/arXiv.2106.00456
59. Li W, Tong J, Anjum MM, Mohammed N, Chen Y, Jiang X. Federated learning algorithms for generalized mixed-effects model (GLMM) on horizontally partitioned data from distributed sources. *BMC Medical Informatics and Decision Making Springer*; 2022;22(1):269. doi:10.1186/s12911-022-02014-1
60. Wedderburn RWM. On the existence and uniqueness of the maximum likelihood estimates for certain generalized linear models. *Biometrika* 1976 Apr;63(1):27–32. doi:10.1093/biomet/63.1.27
61. Van der Vaart AW. *Asymptotic statistics*. Cambridge university press; 2000. ISBN:9780511802256
62. Van Der Vaart AW, Wellner JA. *Weak convergence and empirical processes: With applications to statistics*. Springer; 1996. ISBN:9781475725476


Supplementary Files

Figures

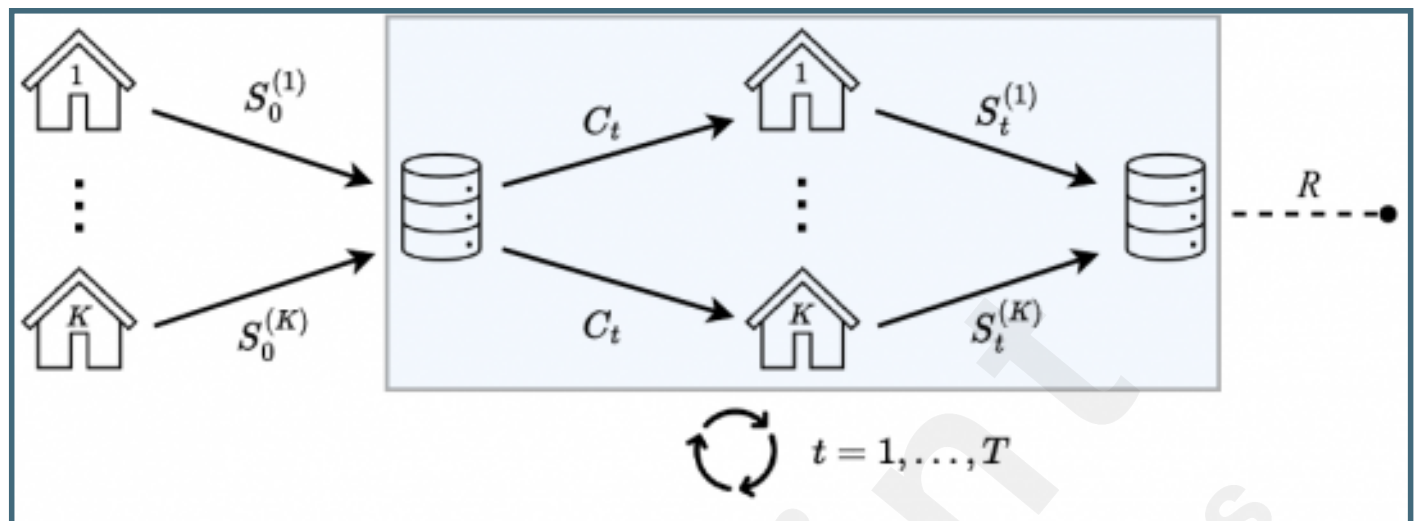
Article selection process for the scoping review. Detailed inclusion and exclusion criteria are described in the text and in the protocol.



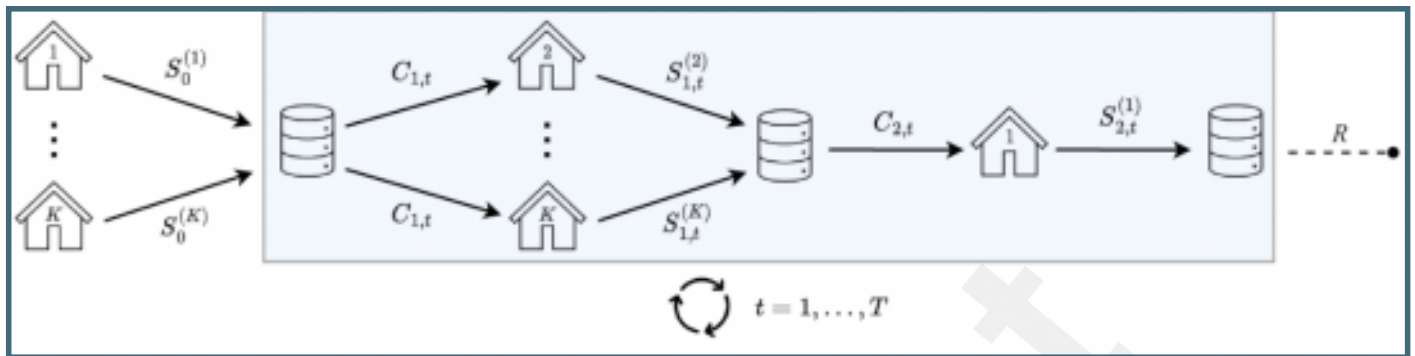
Workflow I: each node calculates summary statistics from its own samples. Results are sent to the CC which combines the information provided by each node to produce the final estimates.



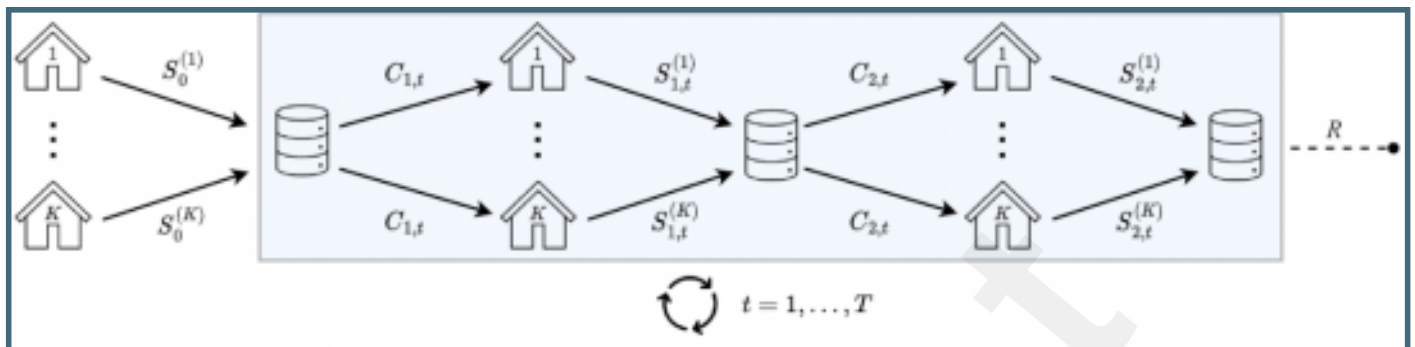
Workflow II: multiple communication rounds are allowed between the CC and the data storage nodes.



Workflow III: multiple communication rounds are allowed between the CC and the data storage nodes, with node 1 following a distinct communication pattern compared to the other nodes.



Workflow IV: multiple communication rounds are allowed between the CC and the data storage nodes, with two back-and-forth distinct communication exchanges between each node and the CC at each iteration.



Multimedia Appendixes

Detailed protocol for the scoping review.

URL: <http://asset.jmir.pub/assets/24940db80f36b41e5d18e8e757dc43d8.docx>

Mathematical derivations pertaining to Objective 3.

URL: <http://asset.jmir.pub/assets/396bea34037bd45bb7d4c95ce1de085d.docx>



CONSORT (or other) checklists

Appendix C PRISMA-ScR checklist.

URL: <http://asset.jmir.pub/assets/b515aae541358e0630fa79820425dfb6.pdf>