Conformal Frequency Estimation with Sketched Data

Anonymous Author(s)
Affiliation
Address
email

Abstract

A flexible conformal inference method is developed to construct confidence intervals for the frequencies of queried objects in very large data sets, based on a much smaller sketch of those data. The approach is data-adaptive and requires no knowledge of the data distribution or of the details of the sketching algorithm; instead, it constructs provably valid frequentist confidence intervals under the sole assumption of data exchangeability. Although our solution is broadly applicable, this paper focuses on applications involving with the count-min sketch algorithm and a non-linear variation thereof. The performance is compared to that of frequentist and Bayesian alternatives through simulations and experiments with data sets of SARS-CoV-2 DNA sequences and classic English literature.

1 Introduction

1.1 Frequency queries from sketched data

Efficiently finding the frequency of an object from a sketch of a very large data set is an important problem in computer science [1, 2], with applications in many fields including machine learning [3], cybersecurity [4], natural language processing [5], genetics [6], and privacy [7]. The count-min sketch (CMS) of [8] is a renowned algorithm for compressing a data set of m objects $Z_1, \ldots, Z_m \in \mathcal{Z}$, belonging to a discrete (and possibly infinite) space $\mathcal{Z}$, into a representation with reduced memory footprint, while allowing simple approximate queries about the observed frequency of any possible $z \in \mathcal{Z}$. At the heart of the CMS lie $d \geq 1$ different w-wide *hash* functions $h_j : \mathcal{Z} \to [w] := \{1, \ldots, w\}$, for all $j \in [d] := \{1, \ldots, d\}$ and some integer $w \geq 1$. Each hash function maps the elements of $\mathcal{Z}$ into one of w possible buckets, and it is designed to ensure that distinct values of z populate the buckets uniformly. Hash functions are typically chosen at random from a *pairwise independent* family $\mathcal{H}$, which ensures the probability (over the randomness in the choice of hash functions) that two distinct objects $z_1, z_2 \in \mathcal{Z}$ are mapped by two different hash functions into the same bucket is equal to $1/w^2$. The data set $Z_1, \ldots, Z_m$ is then compressed into a sketch matrix $C \in \mathbb{N}^{d \times w}$ with row sums equal to m. The element in the j-th row and k-th column of C counts the number of data points mapped by the j-th hash function into the k-th bucket:

$$C_{j,k} = \sum_{i=1}^{m} \mathbb{1}\left[h_j(Z_i) = k\right], \qquad j \in [d]. \tag{1}$$

Because d and w are such that $d \cdot w \ll m$, the matrix C loses information compared to the full data set, but it has the advantage of requiring much less space to store.

Given a sketch C from (1), one may want to answer queries about the original data. A fundamental instance of this problem consists in estimating the empirical frequency of an object $z \in \mathcal{Z}$:

$$f_m(z) := \sum_{i=1}^{m} \mathbb{1}\left[Z_i = z\right]. \tag{2}$$

Submitted to 36th Conference on Neural Information Processing Systems (NeurIPS 2022). Do not distribute.

One solution is to return the smallest count among the d buckets into which z is mapped:

$$\hat{f}_{\text{up}}^{\text{CMS}}(z) = \min_{j \in [d]} \{C_{j,h_j(z)}\} \geq f_m(z). \tag{3}$$

This is a deterministic upper bound for the true $f_m(z)$ [8]. Of course, $\hat{f}_{\text{up}}^{\text{CMS}}(z)$ tends to overestimate $f_m(z)$ due to possible hash collisions. However, the independence of the hash functions ensures that collisions are not too common, leading to a probabilistic lower bound for the true $f_m(z)$.

Theorem 1 ([8])**.** *Fix any $\delta, \epsilon \in (0,1)$. Suppose $d = \lceil -\log \delta \rceil$ and $w = \lceil e/\epsilon \rceil$. Then, $\mathbb{P}_{\mathcal{H}}[f_m(z) \geq \hat{f}_{\text{up}}^{\text{CMS}}(z) - \epsilon m] \geq 1 - \delta$ for any fixed $z \in \mathcal{Z}$, where $\hat{f}_{\text{up}}^{\text{CMS}}(z)$ is as in* (1) *and depends on d, w.*

For example, if $d = 3$, a 95% lower bound on $f_m(z)$ is $\hat{f}_{\text{up}}^{\text{CMS}}(z) - \lceil e \cdot m/w \rceil$. The subscript $\mathcal{H}$ in Theorem 1 clarifies the randomness is with respect to the hash functions, while $Z_1, \ldots, Z_m$ and z are fixed. Although this bound is useful, e.g., to choose d and w prior to sketching, it is not fully satisfactory. First, it is often very conservative if the data are randomly sampled from some distribution as opposed to adversarially fixed. Second, it is not flexible: δ is fixed by d (as opposed to being determined by the user), while ϵ is fixed only by the hash width. A bootstrap method was developed by [9] to address these issues by computing more adaptive lower bounds for $f_m(z)$, under the additional assumptions the data and the query are independent random samples from some underlying distribution. This approach is nearly exact and optimal for the CMS, up to a possible finite-sample discrepancy between the bootstrap and population distributions, but it relies on the specific *linear* structure of the CMS—the sketch matrix C in (3) can be written as a linear combination of the true frequencies of all objects in the data set—and the idea is thus not easily extendable to other sketching algorithms that may perform better than the CMS in practice.

A limitation of the CMS is that it is relatively sensitive to random hash collisions, which can result in overly conservative deterministic upper bounds. Fortunately, a simple modification provides meaningful improvements: the CMS with *conservative updates* (CMS-CU). Whenever a new object Z_i is sketched, the CMS-CU only updates the row of C with the smallest value of $C_{j,h_j(z)}$, leaving the other counters unaltered. Then, the same CMS deterministic upper bound in (3) remains valid. The procedure is outlined in Algorithm 1, Appendix 1. The CMS-CU can greatly improve query accuracy compared to the vanilla CMS [10], but this practical advantage comes at some cost because the algorithm is non-linear. In particular, the theoretical analysis of the CMS-CU beyond a deterministic upper bound is more challenging, and it appears to be a relatively less explored topic.

1.2 Contributions

We develop a novel method for constructing tight and provably valid frequentist confidence intervals for frequency queries based on sketched data, at any desired level. This will work seamlessly with the CMS-CU as well as with any other sketching algorithm, regardless of whether it is linear. Our solution is a non-parametric statistical one, since it makes minimal exchangeability assumptions about the data generating process: its inferences leverage the sampling randomization, not the specific structure of the sketch or the randomness in its hash functions. Although this method requires higher computational power and memory usage compared to traditional sketching, the difference in these requirements will be negligible compared to the size of the data sets involved. As we make use of conformal inference, a brief review of the relevant background is provided in the next section.

1.3 Related work

Different types of lower bounds have been developed for the CMS [11], but most are similar to that of Theorem 1 in the sense that they make no assumptions about the data and only utilize the hashing randomness. An alternative Bayesian perspective is taken by [12, 13]: they model the data with a prior and compute the posterior of the queried frequencies given the sketch. Our approach is closer to that of [9] in that we seek similar frequentist guarantees by leveraging the randomness in the data. However, the theoretical frameworks and methodologies are different: [9] exploits the specific linear structure of the CMS, while we view the sketch as a black box and rely on conformal inference. Conformal inference was pioneered by Vovk and collaborators [14] and brought to the statistics spotlight by [15]. Although primarily conceived for supervised prediction [16–20], conformal inference has found many other applications including outlier detection [21], causal inference [22], and survival analysis [23]. However, its potential for sketching remained untapped. Finally, there exist many

algorithms for using sketches to compute approximate frequency queries; some are similar to the CMS [11, 24–26], while others are more complicated and may involve complex learning algorithms [27, 28]. In theory our method is applicable to all of them, but we focus more on the CMS here because it is a classic algorithm and it provides a familiar example that facilitates the exposition.

2 Preliminaries on conformal prediction

Conformal prediction is typically discussed in a *supervised learning* context, where the data samples are pairs (X_i, Y_i), with X_i indicating a vector of *features* for the i-th observation and Y_i denoting the corresponding continuous (or discrete) *label*. In supervised learning, the goal is to use the information contained in $(X_1, Y_1), \ldots, (X_m, Y_m)$ to learn a mapping between features and labels that can allow one to predict as accurately as possible the unseen label Y_{m+1} of a new sample with features X_{m+1}. In particular, conformal prediction assumes that $(X_1, Y_1), \ldots, (X_{m+1}, Y_{m+1})$ are exchangeable random samples from some unknown joint distribution over (X, Y) and then constructs a prediction interval $[\hat{L}_{m,\alpha}(X_{m+1}), \hat{U}_{m,\alpha}(X_{m+1})]$ with guaranteed *marginal coverage*:

$$\mathbb{P}[\hat{L}_{m,\alpha}(X_{m+1}) \leq Y_{n+1} \leq \hat{U}_{m,\alpha}(X_{m+1})] \geq 1 - \alpha, \tag{4}$$

for any fixed $\alpha \in (0, 1)$. Conformal prediction is flexible, as it can leverage any machine learning algorithm to approximate the true relation between X and Y, thus yielding relatively short intervals satisfying (4). Note that, while it is appropriate to focus on conformal intervals in this paper, similar techniques can also be utilized to construct more general prediction sets [29, 30].

The simplest version of conformal prediction begins by randomly splitting the available observations into two disjoint subsets, assumed for simplicity to have equal size $n = m/2$. The first n samples are spent to fit a black-box machine learning model for predicting Y given X; e.g., a neural network or a random forest. The out-of-sample predictive accuracy of this model is then measured in terms of a *conformity score* for each of the n hold-out data points. In combination with the model learnt from the first half of the data, the empirical distribution of these scores is translated into a recipe for constructing prediction intervals for future test points as a function of X_{m+1}; these intervals are guaranteed to cover Y_{n+1} with probability at least $1 - \alpha$, treating all data as random. The details of this procedure will be clarified shortly. Importantly, the coverage holds exactly in finite samples, regardless of the accuracy of the machine learning model, as long as X_{m+1} is exchangeable with the n hold-out data points. Note that it is unnecessary for the training data to be also exchangeable, as these may be seen as fixed, but it is sometimes easier to say all data points are exchangeable.

The implementation of conformal inference depends on the choice of conformity scores. While several different options exist, an intuitive and general explanation is the following. Imagine that associated with the fitted machine learning model there exists a *nested sequence* [31] of prediction intervals $[\hat{L}_{m,\alpha}(x; t), \hat{U}_{m,\alpha}(x; t)]$ for each possible x. This sequence is indexed by $t \in \mathcal{T} \subseteq \mathbb{R}$, which may be either discrete or continuous, and it is increasing: $\hat{L}_{m,\alpha}(x; t_2) \leq \hat{L}_{m,\alpha}(x; t_1)$ and $\hat{U}_{m,\alpha}(x; t_2) \geq \hat{U}_{m,\alpha}(x; t_1)$ for all $t_2 \geq t_1$. Further, assume there exists $t_\infty \in \mathcal{T}$ such that $\hat{L}_{m,\alpha}(x; t_\infty) \leq Y \leq \hat{U}_{m,\alpha}(x; t_\infty)$ almost surely for all x. For example, one may consider a sequence $\hat{\psi}_m(x) \pm t$, where $\hat{\psi}_m$ is a regression function for a bounded label Y given X learnt by the black-box machine learning model and t plays the role of a predictive standard error. Then, the conformity score for a data point with features $X = x$ and label $Y = y$ can be defined as the smallest index t such that y is contained in the sequence of prediction intervals corresponding to x:

$$E(X, Y) := \min \left\{ t \in \mathcal{T} : Y \in [\hat{L}_{m,\alpha}(x; t), \hat{U}_{m,\alpha}(x; t)] \right\}. \tag{5}$$

The conformal prediction rule is then easily stated. Let $\mathcal{I}^{\text{calib}} \subset \{1, \ldots, m\}$ be the subset of hold-out data points, assumed without loss of generality to have size $n = m/2$. Let $\hat{Q}_{n,1-\alpha}$ be the $\lceil (1 - \alpha)(n + 1) \rceil$ smallest value among the n conformity scores $E(X_i, Y_i)$ evaluated for all $i \in \mathcal{I}^{\text{calib}}$. The conformal prediction interval for a new data point with features X_{m+1} is:

$$\left[\hat{L}_{m,\alpha}(X_{m+1}; \hat{Q}_{n,1-\alpha}), \hat{U}_{m,\alpha}(X_{m+1}; \hat{Q}_{n,1-\alpha}) \right]. \tag{6}$$

Intuitively, this satisfies marginal coverage (4) because Y_{m+1} is outside the interval in (6) if and only if the imaginary score $E(X_{m+1}, Y_{m+1})$ is larger than $\hat{Q}_{n,1-\alpha}$. The rest of the proof is a simple exchangeability argument; see [19] for details, or the proof of Theorem 2 below.

3 Conformalized sketching

3.1 Problem statement

Assume the m data points to be sketched, $Z_1, \ldots, Z_m \in \mathcal{Z}$, are exchangeable random samples from some distribution P_Z on $\mathcal{Z}$. Although P_Z may be arbitrary and unknown, our framework is more restrictive compared to the classical deterministic setting reviewed in Section 1.1, which treats the data as fixed and thus can also handle non-stationary streams or adversarial cases. Yet, there are interesting applications in which data exchangeability can be justified; this will be discussed in more detail in Sections 4 and 5. Consider an arbitrary *sketching* function $\phi : [\mathcal{Z}]^m \to \mathcal{C}$, where $\mathcal{C}$ is some space with lower dimensions compared to $[\mathcal{Z}]^m$. For example, ϕ may represent the sketching performed by the CMS or CMS-CU, in which case $\mathcal{C} = \mathbb{N}^{d \times w}$, for some d, w such that $d \cdot w \ll m$. In general, our framework treats ϕ as a *black box* and allows it to be anything, regardless of whether it involves random hashing or any other specific data compression technique. The goal is to leverage the data exchangeability and the information in $\phi(Z_1, \ldots, Z_m)$ to estimate the unobserved true empirical frequency $f_m(z)$ of some object $z \in \mathcal{Z}$, defined as in (2), while accounting for uncertainty. More precisely, although still informally speaking, we would like to construct a *confidence interval* $[\hat{L}_{m,\alpha}(z), \hat{U}_{m,\alpha}(z)]$ guaranteed to contain $f_m(z)$ "with probability at least $1 - \alpha$", where the randomness here is with respect to the random data sampling.

One way to address this problem is to assume the query z is also randomly sampled exchangeably with $Z_1, \ldots, Z_m$. Thus, we refer to z as Z_{m+1} and focus on computing a confidence interval $[\hat{L}_{m,\alpha}(Z_{m+1}), \hat{U}_{m,\alpha}(Z_{m+1})]$ depending on $\phi(Z_1, \ldots, Z_m)$ that is as short as possible and such that

$$\mathbb{P}\big[\hat{L}_{m,\alpha}(Z_{m+1}) \leq f_m(Z_{m+1}) \leq \hat{U}_{m,\alpha}(Z_{m+1})\big] \geq 1 - \alpha, \tag{7}$$

where the probability is with respect to $Z_1, \ldots, Z_{m+1}$. The interpretation of (7) is as follows: the interval will cover the true frequency $f_m(z)$ for at least fraction $1 - \alpha$ of all future test points on average, if the queries are drawn exchangeably from the same distribution as the sketched data. Of course, this is not fully satisfactory because some objects may be more likely to be queried than others, and therefore not all of our inferences are equally reliable. In particular, the confidence intervals for rare queries may have lower coverage, as illustrated by the following thought experiment. Imagine P_Z has support on $\mathcal{Z} = \{0, 1, 2, \ldots, 10^{100}\}$, with $\mathbb{P}[Z_i = 0] = 0.95$ and $\mathbb{P}[Z_i = z] = 0.05/(|\mathcal{Z}| - 1)$ for all $z \in \mathcal{Z} \setminus \{0\}$ and $i \geq 1$. Then, a 95% confidence interval such as that in (7) may be completely unreliable for all but one specific query about $f_m(0)$. Therefore, we will focus instead on constructing confidence intervals that are simultaneously valid for both rare and random queries, adapting related ideas from the conformalized multi-class classification [14, 29, 32]. Precisely, fix any partition $\mathcal{B} = (B_1, \ldots, B_L)$ of $\{1, \ldots, m\}$ into L bins, such that each B_l is a subinterval of $\{1, \ldots, m\}$, and construct a confidence interval $[\hat{L}_{m,\alpha}(Z_{m+1}), \hat{U}_{m,\alpha}(Z_{m+1})]$ depending on $\phi(Z_1, \ldots, Z_m)$ and $\mathcal{B}$ that is as short as possible and such that:

$$\mathbb{P}\big[\hat{L}_{m,\alpha}(Z_{m+1}) \leq f_m(Z_{m+1}) \leq \hat{U}_{m,\alpha}(Z_{m+1}) \mid f_m(Z_{m+1}) \in B\big] \geq 1 - \alpha, \qquad \forall B \in \mathcal{B}. \tag{8}$$

The choice of $\mathcal{B}$ involves some trade-offs: finer partitions yield more reliable theoretical guarantees but possibly wider confidence intervals; in particular, intervals associated with a bin B may tend to be wider if $\mathbb{P}[Z \in B]$ is small. In practice, we will adopt $|\mathcal{B}| = 5$ later in this paper, but much larger $|\mathcal{B}|$ can be utilized when working with very big data, as it will become clear below.

3.2 Conformalization

In the attempt of adapting conformal inference to address our sketching problem, the first difficulty is that the latter is a data recovery problem, not a supervised prediction task. We propose to overcome this challenge by storing in memory the true frequencies for all objects in the first m_0 observations, with $m_0 \ll m$ but sufficiently large subject to memory constraints. Without loss of generality, assume $m_0 \ll m$; otherwise, the problem would be trivial. Let $n \leq m_0$ indicate the number of distinct objects among the first m_0 observations. The memory required to store these frequencies is $O(n)$, which will typically be negligible as long as m_0 is also small compared to the size of the sketch, $|\mathcal{C}|$. This approach turns our problem into a supervised prediction one, as detailed below.

During an initial *warm-up* phase, the frequencies of the n distinct objects among the first m_0 observations from the data stream, $Z_1, \ldots, Z_{m_0}$, are stored exactly into:

$$f^{\mathrm{wu}}_{m_0}(z) := \sum_{i=0}^{m_0} \mathbb{1}\left[Z_i = z\right]. \tag{9}$$

Next, the remaining $m - m_0$ data points are streamed and sketched, storing also the true frequencies for all instances of objects already seen during the warm-up phase. In other words, the following counters are computed and stored along with the sketch $\phi(Z_{m_0+1}, \ldots, Z_m)$:

$$f^{\mathrm{sv}}_{m-m_0}(z) := \begin{cases} \sum_{i=m_0+1}^{m} \mathbb{1}\left[Z_i = z\right], & \text{if } f^{\mathrm{wu}}_{m_0}(z) > 0, \\ 0, & \text{otherwise.} \end{cases} \tag{10}$$

Again, this requires only $O(n)$ memory. Now, let us define the variable Y_i for all $i \in \{1, \ldots, m_0\} \cup \{m+1\}$ as the true frequency of Z_i among $Z_{m_0+1}, \ldots, Z_m$:

$$Y_i := \sum_{i'=m_0+1}^{m} \mathbb{1}\left[Z_{i'} = Z_i\right]. \tag{11}$$

Note that Y_i is observable for $i \in \{1, \ldots, m_0\}$, in which case $Y_i = f^{\mathrm{sv}}_{m-m_0}(Z_i)$. For a new query, Z_{m+1} is the target of inference—in truth, the target is $f_m(Z_i) = Y_i + f^{\mathrm{wu}}_{m_0}(Z_i)$, but the second term is already known exactly. To complete the connection between sketching and supervised conformal prediction, one still needs to define meaningful features X, and this is where the sketch $\phi(Z_{m_0+1}, \ldots, Z_m)$ comes into play. For each $i \in \{1, \ldots, m_0\} \cup \{m+1, \ldots\}$, define X_i as the vector containing the object of the query as well as all the information in the sketch:

$$X_i := \left(Z_i, \phi(Z_{m_0+1}, \ldots, Z_m)\right). \tag{12}$$

The following result establishes that the pairs $(X_1, Y_1), \ldots, (X_{m_0}, Y_{m_0}), (X_{m+1}, Y_{m+1})$ are exchangeable with one another. This anticipates that conformal prediction techniques can be applied to the supervised observations $(X_1, Y_1), \ldots, (X_{m_0}, Y_{m_0})$ to predict Y_{m+1} given X_{m+1}, guaranteeing the marginal coverage property in (7). All mathematical proofs are in Appendix 2.

Proposition 1. *If the data $Z_1, \ldots, Z_{m+1}$ are exchangeable, then the pairs of random variables $(X_1, Y_1), \ldots, (X_{m_0}, Y_{m_0}), (X_{m+1}, Y_{m+1})$ defined in* (11)–(12) *are exchangeable with one another.*

Confidence intervals satisfying the stronger conditional coverage property in (8) can be obtained by modifying the standard conformal inference procedure as follows. First, each data point (X_i, Y_i) in the calibration set $\mathcal{I}^{\mathrm{calib}} = \{1, \ldots, m_0\}$ is assigned to the appropriate frequency bin based on Y_i. Define n_l as the number of calibration data points assigned to bin B_l, for all $l \in \{1, \ldots, L\}$. Then, $\hat{Q}_{n_l, 1-\alpha}(B_l)$ is defined as the $\lceil (1-\alpha)(n_l+1) \rceil$ smallest value among the n_l conformity scores assigned to bin B_l. Finally, the confidence interval for a new random query X_{m+1} is given by

$$\left[\hat{L}_{m,\alpha}(X_{m+1}; \hat{Q}^*_{n,1-\alpha}), \hat{U}_{m,\alpha}(X_{m+1}; \hat{Q}^*_{n,1-\alpha})\right],$$

where

$$\hat{Q}^*_{n,1-\alpha} := \max_{l \in \{1,\ldots,L\}} \hat{Q}_{n_l, 1-\alpha}(B_l),$$

and $[\hat{L}_{m,\alpha}(\cdot; t), \hat{U}_{m,\alpha}(\cdot; t)]$ is a rule for computing a nested sequence of intervals depending on Z_{m+1} and $\phi(Z_{m_0+1}, \ldots, Z_m)$. Examples of such rules are in the next section. Importantly, these rules may involve parameters to be fitted on a subset of m_0^{train} supervised data points (X_i, Y_i) for $i \in \{1, \ldots, m_0^{\mathrm{train}}\}$, as long as the scores are only evaluated on the remaining $m_0 - m_0^{\mathrm{train}}$ points. This procedure is outlined in Algorithm 2, Appendix 1. The following result states that the confidence interval output by Algorithm 2 has the desired frequency-range conditional coverage.

Theorem 2. *If the data $Z_1, \ldots, Z_{m+1}$ are exchangeable, the confidence interval output by Algorithm 2 satisfies the frequency-range conditional coverage property defined in* (8).

3.3 Conformity scores

Algorithm 2 can accommodate virtually any data-adaptive rule, depending on Z_{m+1} and on the sketch $\phi := \phi(Z_{m_0+1}, \ldots, Z_m)$, for computing nested confidence intervals. Two concrete options are presented here. For simplicity, we focus on one-sided confidence intervals; that is, we seek a $1-\alpha$ lower bound for $f_m(X_{m+1})$. This is useful when a deterministic upper bound $\hat{f}_{\mathrm{up}}(Z_{m+1}) \geq f_m(Z_{m+1})$ is already available, as it is the case with the CMS or CMS-CU. Then, one can simply set $\hat{U}_{m,\alpha}((z,\phi);t) := \hat{f}_{\mathrm{up}}(z)$ and focus on computing $\hat{L}_{m,\alpha}(\cdot;t)$.

The first option for $\hat{L}_{m,\alpha}(\cdot;t)$ is to use the following fixed rule:

$$\hat{L}^{\mathrm{fixed}}_{m,\alpha}((z,\phi);t) := \max\{0, \hat{f}_{\mathrm{up}}(Z_{m+1}) - t\}, \qquad t \in \{0,1,\ldots,m\}. \tag{13}$$

In words, the lower bound for $f_m(Z_{m+1})$ in (14) is defined by shifting the deterministic upper bound downward by t. The appropriate value of t guaranteeing the desired coverage for future random queries is calculated by the conformal inference procedure outlined in Algorithm 2. This approach has the advantage of being simple and intuitive, and it is very similar to the optimal solution of [9] for the special case of the CMS. Further, it does not need training data, so all m_0 observations with tracked frequencies can be utilized for computing conformity scores.

The second option involves training but has the advantage of being more flexible; this is inspired by the methods of [33, 34] for regression. Concretely, consider a machine learning model that takes as input the known upper bound $\hat{f}_{\mathrm{up}}(Z_i)$ and estimates the conditional distribution of $\hat{f}_{\mathrm{up}}(Z_i) - f_m(Z_i)$ given $\hat{f}_{\mathrm{up}}(Z_i)$. For example, think of a multiple quantile neural network [35] or a quantile random forest [36]. After fitting this model on the m_0^{train} supervised data points (X_i, Y_i) allocated for training, let $\hat{q}_t$ be the estimated α_t lower quantile of the of $\hat{f}_{\mathrm{up}}(Z_i) - f_m(Z_i) \mid \hat{f}_{\mathrm{up}}(Z_i)$, for all $t \in [1,\ldots,T]$ and some fixed sequence $1 = \alpha_1 < \ldots < \alpha_T = 0$. Without loss of generality, assume the machine learning model is that $\hat{q}_0 = 0$ and $\hat{q}_T = m$. Then, define $\hat{L}_{m,\alpha}(\circ;t)$ as:

$$\hat{L}^{\mathrm{adaptive}}_{m,\alpha}((z,\phi);t) := \max\left\{0, \hat{f}_{\mathrm{up}}(X_{m+1}) - \hat{q}_t\left(\hat{f}_{\mathrm{up}}(X_{m+1})\right)\right\}, \qquad t \in \{0,1,\ldots,m\}. \tag{14}$$

This second approach can lead to a lower bound whose distance from the upper bound is data adaptive. This is an advantage because some sketching algorithms may introduce higher uncertainty about common queries compared to rarer ones, or vice versa, and such patterns can be learnt given sufficient data. Of course, the two above examples of $\hat{L}_{m,\alpha}(\cdot;t)$ are not exhaustive. Algorithm 2 can be applied in combination with any rule for computing nested sequences of lower bounds, and it can leverage all high-dimensional information contained in $\phi(Z_{m_0+1}, \ldots, Z_m)$, not just the deterministic upper bound for the CMS and CMS-CU. However, the applications in this paper focus on the CMS and CMS-CU, so other families of lower bounds are not discussed here.

4 Applications

4.1 Experiments with synthetic data sets

Conformalized sketching is applied with the two types of conformity scores described in Section 3.3. The adaptive scores are based on an isotonic distributional regression model [37]. The goal is to compute lower frequency bounds for random queries based on simulated data compressed by the CMS-CU, implemented with $d = 3$ hash functions of width $w = 1000$. In particular, $m = 100,000$ observations are sampled i.i.d. from some distribution specified below. The first $m_0 = 5000$ observations are stored without loss during the warm-up phase, as outlined in Algorithm 2, while the remaining $95,000$ are compressed by the CMS-CU. The conformity scores are evaluated separately within $L = 5$ frequency bins, seeking the frequency-range conditional coverage property defined in (8). The bins are determined in a data-driven fashion so that each contains approximately the same probability mass; in practice, this is achieved by partitioning the range of frequencies for the objects tracked exactly by Algorithm 2 according to the observed empirical quantiles. Lower bounds for new queries are computed for $10,000$ data points also sampled i.i.d. from the same distribution. The quality of these bounds is quantified with two metrics: the mean *length* of the resulting confidence intervals and the coverage the proportion of queries for which the true frequency is correctly covered, or empirical *coverage*. The performance is averaged over 10 independent experiments.

Experiments are performed on synthetic data sampled from two families of distributions. First, we consider a Zipf distribution with $\mathbb{P}[Z_i = z] = z^{-a}/\zeta(a)$ for all $z \in \{1, 2 \ldots, \}$, where ζ is the Riemann Zeta function and $a > 1$ is a control parameter affecting the power-law tail behavior. Second, synthetic data are generated from a random probability measure distributed as the Pitman-Yor prior [38] with a standard Gaussian base distribution and parameters $\lambda > 0$ and $\sigma \in [0, 1)$, as explained in Appendix 1.3. The parameter λ is set to $\lambda = 5000$, while σ is varied. For $\sigma = 0$ the Pitman-Yor prior reduces to the Dirichlet prior [39], while $\sigma > 0$ results in heavier tails.

Three benchmark methods are considered. The first one is the *classical* 95% lower bound in Theorem 1. The second one is the *Bayesian* method of [12], which assumes a non-parametric Dirichlet process prior for the distribution of the data stream, estimates its scaling parameter by maximizing the marginal likelihood of the observed sketch, and then computes the posterior of the queried frequencies conditional on the observed sketch. The performance of the lower 5% posterior quantile is compared to the lower bounds obtained with the other methods according to the frequentist metrics defined above. The third benchmark is the bootstrap method of [9], which is nearly exact and optimal for the vanilla CMS (up to some possible finite-sample discrepancy between the bootstrap and population distributions) but is not theoretically valid for other sketching techniques.

Figure 1 compares the performance of the proposed conformal method to those of the three benchmarks on the Zipf data. All methods achieve marginal coverage (4), with the exception of the Bayesian approach which in this case is based on a misspecified prior. The length of the confidence intervals indicates the classical bound is very conservative, while the bootstrap and conformal methods provide relative informative bounds, particularly when a is larger and hash collisions become rarer. The conformal intervals can be the shortest ones, especially if implemented with the adaptive conformity scores. This should not be surprising because the bootstrap approach is not guaranteed to be optimal for the non-linear compression induced by the CMS-CU. Indeed, as shown in Figure 2, the machine learning model deployed by our adaptive scores can detect and take advantage of the fact that this non-linear sketch allows increased precision for more common queries.

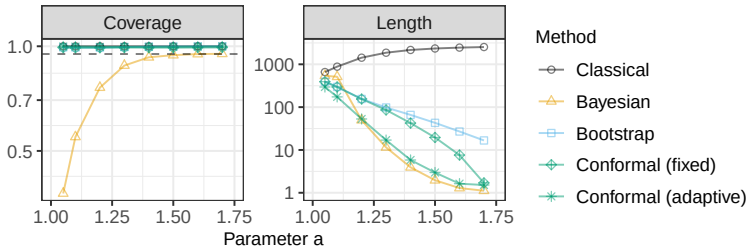

Figure 1: Coverage and length of 95% confidence intervals with data from a Zipf distribution, sketched with CMS-CU. The results are shown as a function of the Zipf tail parameter a. Standard errors would be too mall to be clearly visible in this figure, and are hence omitted.

Supplementary results reported in Appendix 3 show the CMS-CU leads to more precise queries with all methods compared to the vanilla CMS; see Figure A1. Figure A2 confirms that conformal lower bounds no longer have a clear advantage over the bootstrap ones if the data are sketched with the vanilla CMS instead of the CMS-CU. In fact, although the conformal intervals obtained with the adaptive conformity scores can be little shorter than the bootstrap ones even for the vanilla CMS, the latter have the advantage of (approximately, in the limit of large samples) satisfying an even stronger frequency-conditional coverage property equivalent to (8) with bins of size 1 [9]. Analogous results for the experiments with Pitman-Yor process data are also in Appendix 3; see Figures A3–A6.

4.2 Analysis of 16-mers in SARS-CoV-2 DNA sequences

This application involves a data set of nucleotide sequences from SARS-CoV-2 viruses made publicly available by the National Center for Biotechnology Information [40]. The data include 43,196 sequences, each consisting of approximately 30,000 nucleotides. The goal is to estimate the empirical frequency of each possible *16-mer*, a distinct sequence of 16 DNA bases in contiguous nucleotides. Given that each nucleotide has one of 4 bases, there are $4^{16} \approx 4.3$ billion possible 16-mers.

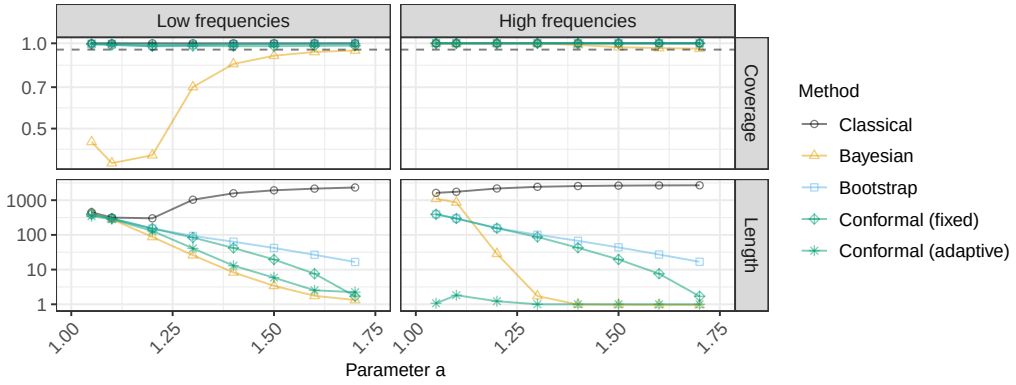

Figure 2: Performance of confidence intervals stratified by the true query frequency. Left: frequency below median; right: frequency above median. Other details are as in Figure 1.

Thus, exact tracking of all 16-mers is not unfeasible, which allows us to validate the sketch-based queries. Sequences containing missing values are removed during pre-processing, for simplicity.

The experiments are carried out as in Section 4.1, but a larger sample of 1,000,000 observations is sketched, and the width w of the hash functions is varied as a control parameter. Figure 3 compares the performances of all methods as a function of the hash width, measured in terms of marginal coverage and mean width of the confidence intervals. All methods appear to have marginal coverage above the nominal level, except for the Bayesian approach which yields slightly lower coverage if w is large. For small hash widths, all methods construct confidence intervals of similar width, because the distribution of SARS-CoV-2 16-mers frequencies is quite concentrated with relatively narrow support (see Figure A7 in Appendix 3), which makes it especially difficult to compress the data without much loss. By contrast, the proposed conformal methods yield noticeably shorter confidence intervals if w is large. Figure A8 reports the same results stratified by the frequency of the queried objects, while Figure A9 confirms the advantage of sketching with the CMS-CU as opposed to the vanilla CMS. Table A1 visualizes a list of 10 common and 10 rare queries along with their corresponding deterministic upper bounds for $w = 50,000$, comparing the lower bounds obtained with each method. Table A2 shows analogous results obtained with $w = 5,000$.

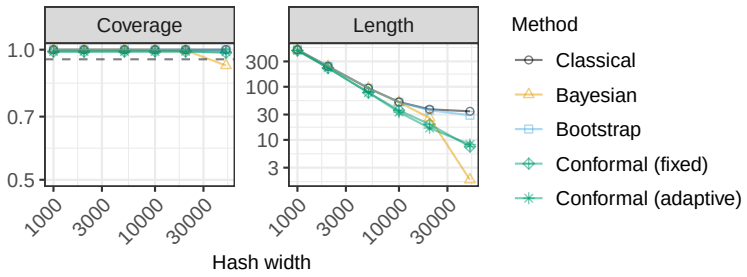

Figure 3: Performance of confidence intervals based on sketched SARS-CoV-2 sequence data. The results are shown as a function of the hash width. Other details are as in Figure 1.

Figure A10 compares the performances of different frequency *point-estimates* measured in terms of mean absolute deviation from the true frequency. With the classical method, we take for simplicity the midpoint of the 95% confidence interval as a point estimate, although other approaches are also possible [11]. For the other methods, the point estimate is defined as the lower confidence bound at level $\alpha = 0.5$; in the Bayesian case, this is the posterior median. Although a conformal lower bound at level $\alpha = 0.5$ is not always a reliable estimator of conditional medians [41], because the conformal coverage guarantees treat the query as random, this approach consistently outperformed the benchmarks in all of our experiments.

4.3 Analysis of 2-grams in English literature

This application is based on a data set consisting of 18 open-domain classic pieces of English literature downloaded using the NLTK Python package [42] from the Gutenberg Corpus [43]. The goal is to count the frequencies of all *2-grams*—consecutive pairs of English words—across this corpus. After some basic preproccessing to remove punctuation and unusual words (only those in a dictionary of size 25,487 common English words are retained), the total number of 2-grams left in this data set is approximately 1,700,000 (although the total number of all *possible* 2-grams within this dictionary is approximately 650,000,000). The same experiments are then carried out as in Section 4.2, sketching 1,000,000 randomly sampled 2-grams and querying 10,000 independent 2-grams.

Figure 4 shows the conformal intervals with adaptive scores achieve the desired coverage and tend to have the shortest width, while the Bayesian intervals are not valid unless the hashes are very wide. The conformal approach has a larger advantage here because these data can be compressed more efficiently compared to the one in the previous section because the frequency distribution of English 2-grams has power-law tails; see Figure A7. Additional results along the lines of those in the previous section are in Appendix 3; see Figures A11–A13 and Tables A1–A2.

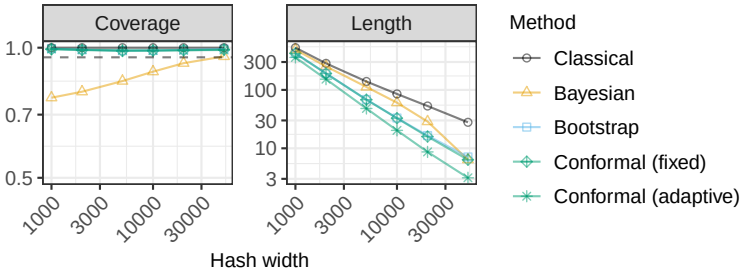

Figure 4: Performance of confidence intervals for random frequency queries, for a sketched data set of English 2-grams in classic English literature. Other details are as in Figure 3.

5 Discussion

Conformalized sketching is a non-parametric and data-adaptive statistical method for quantifying uncertainty in problems involving frequency estimation from sketched data. This paper has revolved around the CMS because that is a prominent technique for which several benchmarks are available. However, our method is very broadly applicable because it operates under the sole assumption of data exchangeability, without requiring any knowledge of the sketching algorithm. Of course, exchangeability is not always an appropriate assumption, and thus one may sometimes need to rely on more conservative bounds based on hashing randomness. Yet there are interesting situations in which the sketched data can be seen as exchangeable random samples from a population. For example, in natural language processing one may wish to count the frequencies of non-contiguous tuples of words co-occurring in the same sentence within a large corpus, expanding on the example of Section 4.3. In that case, it would be unfeasible to keep track of all frequencies exactly, but sketching random samples would be relatively easy. One limitation of conformalized sketching is that it does not provide guarantees about the expected proportion of correct *unique* queries. In fact, if the queries are randomly sampled exchangeably with the sketched data points, some of them may be redundant. Even though the expected proportion of incorrect unique queries is sometimes below α (Figures A14–A15), this is not always the case (Figure A16). It may be possible to modify the procedure to achieve this additional guarantee, but we leave this intriguing problem for future work.

Software and computations

Code and data needed to reproduce these experiments are available in the supplementary material. Experiments were carried out in parallel using a computing cluster; each experiment required less than a few hours of time with a standard CPU and less than 5GB of memory (approximately 20 GB of memory are needed for the analysis of the SARS-CoV-2 DNA data) .

References

[1] Jayadev Misra and David Gries. Finding repeated elements. *Science of computer programming*, 2(2):143–152, 1982.

[2] Moses Charikar, Kevin Chen, and Martin Farach-Colton. Finding frequent items in data streams. In *International Colloquium on Automata, Languages, and Programming*, pages 693–703. Springer, 2002.

[3] Qinfeng Shi, James Petterson, Gideon Dror, John Langford, Alex Smola, and SVN Vishwanathan. Hash kernels for structured data. *Journal of Machine Learning Research*, 10(11), 2009.

[4] Stuart Schechter, Cormac Herley, and Michael Mitzenmacher. Popularity is everything: A new approach to protecting passwords from {Statistical-Guessing} attacks. In *5th USENIX Workshop on Hot Topics in Security (HotSec 10)*, 2010.

[5] Amit Goyal, Hal Daumé III, and Graham Cormode. Sketch algorithms for estimating point queries in nlp. In *Proceedings of the 2012 joint conference on empirical methods in natural language processing and computational natural language learning*, pages 1093–1103, 2012.

[6] Qingpeng Zhang, Jason Pell, Rosangela Canino-Koning, Adina Chuang Howe, and C Titus Brown. These are not the k-mers you are looking for: efficient online k-mer counting using a probabilistic data structure. *PloS one*, 9(7):e101271, 2014.

[7] Graham Cormode, Somesh Jha, Tejas Kulkarni, Ninghui Li, Divesh Srivastava, and Tianhao Wang. Privacy at scale: Local differential privacy in practice. In *Proceedings of the 2018 International Conference on Management of Data*, pages 1655–1658, 2018.

[8] Graham Cormode and Shan Muthukrishnan. An improved data stream summary: the count-min sketch and its applications. *Journal of Algorithms*, 55(1):58–75, 2005.

[9] Daniel Ting. Count-min: Optimal estimation and tight error bounds using empirical error distributions. In *Proceedings of the 24th ACM SIGKDD International Conference on Knowledge Discovery & Data Mining*, pages 2319–2328, 2018.

[10] Cristian Estan and George Varghese. New directions in traffic measurement and accounting. In *Proceedings of the 2002 conference on Applications, technologies, architectures, and protocols for computer communications*, pages 323–336, 2002.

[11] Graham Cormode and Ke Yi. *Small summaries for big data*. Cambridge University Press, 2020.

[12] Diana Cai, Michael Mitzenmacher, and Ryan P Adams. A Bayesian nonparametric view on count-min sketch. In *Advances in Neural Information Processing Systems 31*, pages 8782–8791, 2018.

[13] Emanuele Dolera, Stefano Favaro, and Stefano Peluchetti. A Bayesian nonparametric approach to count-min sketch under power-law data streams. In *International Conference on Artificial Intelligence and Statistics*, pages 226–234. PMLR, 2021.

[14] Vladimir Vovk, Alex Gammerman, and Glenn Shafer. *Algorithmic learning in a random world*. Springer, 2005.

[15] Jing Lei, Max G'Sell, Alessandro Rinaldo, Ryan J. Tibshirani, and Larry Wasserman. Distribution-free predictive inference for regression. *Journal of the American Statistical Association*, 113(523):1094–1111, 2018.

[16] Vladimir Vovk, Ilia Nouretdinov, and Alex Gammerman. On-line predictive linear regression. *Annals of Statistics*, 37(3):1566–1590, 2009.

[17] Vladimir Vovk. Cross-conformal predictors. *Annals of Mathematics and Artificial Intelligence*, 74(1-2):9–28, 2015.

[18] Jing Lei and Larry Wasserman. Distribution-free prediction bands for non-parametric regression. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)*, 76(1):71–96, 2014.

[19] Yaniv Romano, Evan Patterson, and Emmanuel Candès. Conformalized quantile regression. In *Advances in Neural Information Processing Systems*, pages 3538–3548, 2019.

[20] Rafael Izbicki, Gilson T Shimizu, and Rafael B Stern. Flexible distribution-free conditional predictive bands using density estimators. *preprint at arXiv:1910.05575*, 2019.

[21] Stephen Bates, Emmanuel Candès, Lihua Lei, Yaniv Romano, and Matteo Sesia. Testing for outliers with conformal p-values. *preprint at arXiv:2104.08279*, 2021.

[22] Lihua Lei and Emmanuel Candès. Conformal inference of counterfactuals and individual treatment effects. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)*.

[23] Emmanuel Candès, Lihua Lei, and Zhimei Ren. Conformalized survival analysis. *arXiv preprint arXiv:2103.09763*, 2021.

[24] Li Fan, Pei Cao, Jussara Almeida, and Andrei Z Broder. Summary cache: a scalable wide-area web cache sharing protocol. *IEEE/ACM transactions on networking*, 8(3):281–293, 2000.

[25] Amit Goyal and Hal Daumé. Lossy conservative update (lcu) sketch: Succinct approximate count storage. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 25, pages 878–883, 2011.

[26] Guillaume Pitel and Geoffroy Fouquier. Count-min-log sketch: Approximately counting with approximate counters. *arXiv preprint arXiv:1502.04885*, 2015.

[27] Chen-Yu Hsu, Piotr Indyk, Dina Katabi, and Ali Vakilian. Learning-based frequency estimation algorithms. In *International Conference on Learning Representations*, 2019.

[28] Tanqiu Jiang, Yi Li, Honghao Lin, Yisong Ruan, and David P Woodruff. Learning-augmented data stream algorithms. In *International Conference on Learning Representations*, 2019.

[29] Yaniv Romano, Matteo Sesia, and Emmanuel Candès. Classification with valid and adaptive coverage. *Advances in Neural Information Processing Systems*, 33:3581–3591, 2020.

[30] Anastasios Angelopoulos, Stephen Bates, Jitendra Malik, and Michael I Jordan. Uncertainty sets for image classifiers using conformal prediction. *arXiv preprint arXiv:2009.14193*, 2020.

[31] Chirag Gupta, Arun K Kuchibhotla, and Aaditya K Ramdas. Nested conformal prediction and quantile out-of-bag ensemble methods. *preprint at arXiv:1910.10562*, 2019.

[32] Mauricio Sadinle, Jing Lei, and Larry Wasserman. Least ambiguous set-valued classifiers with bounded error levels. *Journal of the American Statistical Association*, 114(525):223–234, 2019.

[33] Victor Chernozhukov, Kaspar Wüthrich, and Yinchu Zhu. Distributional conformal prediction. *Proceedings of the National Academy of Sciences*, 118(48), 2021.

[34] Matteo Sesia and Yaniv Romano. Conformal prediction using conditional histograms. *Advances in Neural Information Processing Systems*, 34, 2021.

[35] James W. Taylor. A quantile regression neural network approach to estimating the conditional density of multiperiod returns. *Journal of Forecasting*, 19(4):299–311, 2000.

[36] Nicolai Meinshausen. Quantile regression forests. *Journal of Machine Learning Research*, 7: 983–999, 2006.

[37] Alexander Henzi, Johanna F. Ziegel, and Tilmann Gneiting. Isotonic distributional regression. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)*, 83(5):963–993, 2021.

[38] Jim Pitman and Marc Yor. The two parameter poisson-dirichlet distribution derived from a stable subordinator. *The Annals of Probability*, 25:855–900, 1997.

[39] S. Thomas Ferguson. A Bayesian analysis of some nonparametric problems. *The Annals of Statistics*, 1:209–230, 1973.

[40] Eneida L Hatcher, Sergey A Zhdanov, Yiming Bao, Olga Blinkova, Eric P Nawrocki, Yuri Ostapchuck, Alejandro A Schäffer, and J Rodney Brister. Virus variation resource–improved response to emergent viral outbreaks. *Nucleic acids research*, 45(D1):D482–D490, 2017.

[41] Dhruv Medarametla and Emmanuel Candès. Distribution-free conditional median inference. *Electronic Journal of Statistics*, 15(2):4625–4658, 2021.

[42] Steven Bird, Ewan Klein, and Edward Loper. *Natural language processing with Python: analyzing text with the natural language toolkit*. " O'Reilly Media, Inc.", 2009.

[43] Project Gutenberg. Project Gutenberg. `www.gutenberg.org`. Accessed: 2022-02-05.

Checklist

1. For all authors...
   (a) Do the main claims made in the abstract and introduction accurately reflect the paper's contributions and scope? [Yes]
   (b) Did you describe the limitations of your work? [Yes]
   (c) Did you discuss any potential negative societal impacts of your work? [N/A]
   (d) Have you read the ethics review guidelines and ensured that your paper conforms to them? [Yes]
2. If you are including theoretical results...
   (a) Did you state the full set of assumptions of all theoretical results? [Yes]
   (b) Did you include complete proofs of all theoretical results? [Yes]
3. If you ran experiments...
   (a) Did you include the code, data, and instructions needed to reproduce the main experimental results (either in the supplemental material or as a URL)? [Yes]
   (b) Did you specify all the training details (e.g., data splits, hyperparameters, how they were chosen)? [Yes]
   (c) Did you report error bars (e.g., with respect to the random seed after running experiments multiple times)? [Yes]
   (d) Did you include the total amount of compute and the type of resources used (e.g., type of GPUs, internal cluster, or cloud provider)? [Yes]
4. If you are using existing assets (e.g., code, data, models) or curating/releasing new assets...
   (a) If your work uses existing assets, did you cite the creators? [N/A]
   (b) Did you mention the license of the assets? [N/A]
   (c) Did you include any new assets either in the supplemental material or as a URL? [N/A]
   (d) Did you discuss whether and how consent was obtained from people whose data you're using/curating? [N/A]
   (e) Did you discuss whether the data you are using/curating contains personally identifiable information or offensive content? [N/A]
5. If you used crowdsourcing or conducted research with human subjects...
   (a) Did you include the full text of instructions given to participants and screenshots, if applicable? [N/A]
   (b) Did you describe any potential participant risks, with links to Institutional Review Board (IRB) approvals, if applicable? [N/A]
   (c) Did you include the estimated hourly wage paid to participants and the total amount spent on participant compensation? [N/A]