



Privacy-Preserving Machine Learning on Web Browsing for Public Opinion

Sam Buxbaum^{1(✉)}, Lucas M. Tassis¹, Lucas Boschelli², Giovanni Comarella³,
Mayank Varia¹, Mark Crovella¹, and Dino P. Christenson⁴

¹ Boston University, Boston, USA
{sambux, ltassis, varia, crovella}@bu.edu

² Maritz, Fenton, USA
boschellil@wustl.edu

³ Universidade Federal do Espirito Santo, Vitoria, Brazil
gc@inf.ufes.br

⁴ Washington University, St. Louis, USA
dinopc@wustl.edu

Abstract. We present a real-world deployment of secure multiparty computation to predict political preference from private web browsing data. To estimate aggregate preferences for the 2024 U.S. presidential election candidates, we collect and analyze secret-shared data from nearly 8000 users from August 2024 through February 2025, with over 2000 daily active users sustained throughout the bulk of the survey. The use of MPC allows us to compute over sensitive web browsing data that users would otherwise be more hesitant to provide. We collect data using a custom-built Chrome browser extension and perform our analysis using the CrypTen MPC library. To our knowledge, we provide the first implementation under MPC of a model for the learning from label proportions (LLP) problem in machine learning, which allows us to train on unlabeled web browsing data using publicly available polling and election results as the ground truth.

1 Introduction

Secure multi-party computation (MPC) is a cryptographic protocol that allows several people to contribute their data toward a collective data analysis without ever exposing their personal data to any other party. MPC has been a topic of research for decades [10, 29, 62, 71], and through a variety of algorithmic improvements and software implementations (e.g., [13, 39, 47, 51, 63]), it has seen deployments over the past decade in the commercial and public sectors (e.g., [1, 6, 11, 14, 24, 44, 58]). Recently, there has been a focus on specialized MPC algorithms for machine learning operations like gradient descent and logistic regression (e.g., [4, 28, 49, 50, 52]) with corresponding software implementations like CrypTen [40].

In this work, we develop and deploy privacy-preserving machine learning in a real-world application based on political science: namely, the estimation of political preferences in the United States during the months around the 2024 U.S.

presidential election. To accomplish this goal, we combine MPC-based privacy-preserving machine learning with the work of Comarela et al. [18], which demonstrates that web browsing patterns can be used to assess aggregate preferences in political candidates.

Background. In more detail, Comarela et al. [18] start from a dataset of web browsing records of a cross-section of people provided by a media measurement company. They design a machine learning algorithm based on Learning from Label Proportions (LLP), shown in Algorithm 1, that uses web browsing data to infer the same type of information that is generally produced by political polls: the fraction of voters in each region (e.g., county or state) that prefer a given candidate at a particular point in time. Hence, this work shows potential to augment traditional methods of political polling, since its lower cost enables more frequent and precise polling—an opinion poll can form the initial “ground truth” in one location at one moment in time, and then LLP on web browsing data can be used to predict political preferences at other times and/or locations.

This LLP algorithm is remarkably effective at estimating political preferences, even when given only browsing visits to the most popular websites rather than the long tail of smaller blogs and personal sites. Nevertheless, data privacy concerns make it challenging and risky to collect and store sensitive data about web browsing activity in practice.

Our Contributions. In this work, we design, implement, and deploy a privacy-preserving system that assesses candidate preferences while also providing the cryptographic guarantee that web browsing data never leaves client browsers in the clear. In more detail, we contribute the following.

- An MPC system that comprises a client-side browser extension that calculates a daily histogram of visits to popular websites, and a cloud-based MPC backend built on top of CrypTen [40] that uses these secret-shared histograms to calculate aggregate political preferences. We detail the system architectural design in Sect. 2 and describe our algorithmic innovations in Sect. 3.
- A case study of the deployment of our MPC system on a panel of nearly 8000 people during the months around the 2024 U.S. presidential election. We describe our work in Sect. 4 and emphasize upfront that we focus in this paper more on the computational decisions and security considerations involved in deploying MPC in this setting rather than the political science insights resulting from the data analysis.

Ethics. We conclude this section with a brief but important discussion of the ethics of this type of research. We caution that the use of MPC cryptographic technology does not automatically ensure adherence to social, ethical, or legal obligations to protect privacy [68], and as a result we have taken multiple steps to ensure that our deployment respects personal autonomy and social norms of privacy. First, we received approval from an Institutional Review Board to deploy this system, and we provided clear and simple methods for panelists to

learn about their data privacy rights and to opt-out of participation at any time if desired. Second, we made sure to design our privacy-preserving machine learning algorithms to reveal only the overall fraction of people in a region (e.g., a U.S. state) that have a particular political preference, and not to infer or reveal any information about which people have that preference. Third, we designed the client browser extension so that it becomes “dormant” after the end of the study period: in other words, it automatically stops observing and transmitting secret-shared web browsing data even if the panelist never deletes the extension. Fourth, we provided fair compensation to panelists for their participation (as approved by the IRB), and we partitioned our own research team so that the researchers who connected with the panelists and observed their payment information were unable to access any other part of the system, and vice-versa (Fig. 1).

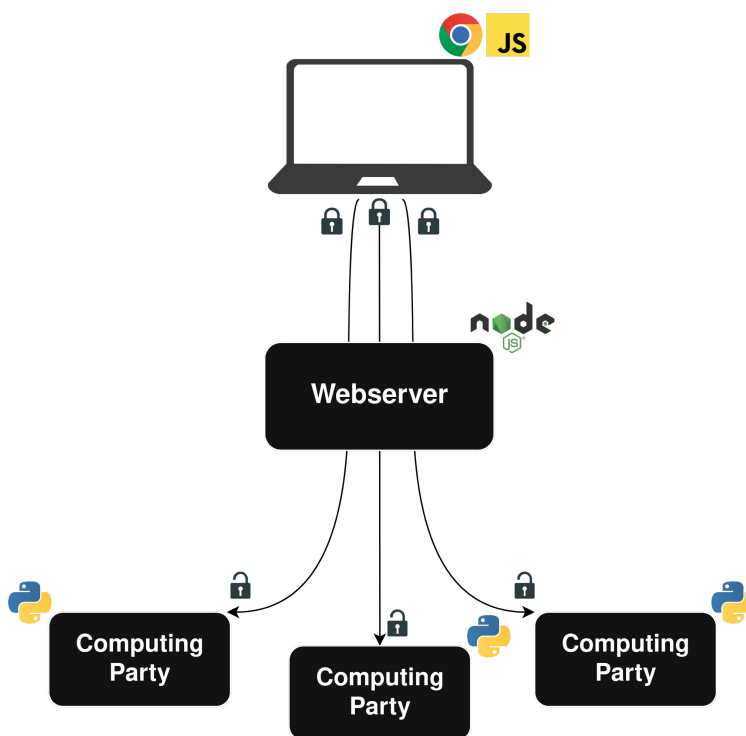


Fig. 1. The OPPS system design. The tools used for each component are included in the upper corner of the component.

2 System Design and Software Architecture

We perform our analysis using MPC in the honest-majority three-party setting with semi-honest security. We operate in the outsourced MPC setting, where users upload secret shares to a set of non-colluding servers who carry out the computation. Concretely, our system contains three components: (i) a browser

extension used by the clients, (ii) a webserver to handle client interaction, and (iii) a set of three non-colluding compute servers to carry out the MPC.

2.1 Browser Extension

We develop a custom Chrome browser extension to monitor users' web browsing behavior. Emulating the ideas in Comarela et al. [18], the extension tracks (i) how many times each user visits each of the most popular 517 websites and (ii) how many times they were referred to each of the top 517 websites from popular social media websites. See Sect. 4.1 for details on how we obtained the list of websites.

The visit and referral histograms are secret-shared in the browser on the client side, and the shares are encrypted under the public keys of the three MPC servers. We implement secret-sharing and hybrid public-key encryption (with RSA and AES) using the JavaScript Web Crypto API. Once per day, the plugin secret shares, encrypts, and sends the histograms to the webserver.

2.2 Webserver

We use a continuously running webserver to handle all communication with the clients and to eliminate the need to have the more expensive computing servers active at all times. The webserver, written in NodeJS, is lightweight and can process many concurrent connections while running on a small machine in the cloud. When a client uploads their encrypted secret shares, the webserver stores them to be retrieved by the computing servers. We emphasize that the webserver exists to simplify the interaction with the clients, but the secret shares are end-to-end encrypted from the client to the computing parties, so the webserver itself learns nothing.

2.3 Computing Servers

We implement the MPC data analysis on top of the CrypTen Python library [40]. The three MPC servers retrieve and decrypt their individual secret shares of the web browsing data from the webserver, and they collectively (and obviously) train a model or perform inference with a previously trained model.

In order to use CrypTen out of the box, all three MPC servers must be accessible by a single AWS account, which presents an obvious privacy challenge. To overcome the challenge, we place all three servers in the same AWS account, and the three server operators are collectively locked out of the account. Computation is performed using an automated process that allows server operators to communicate with the servers without giving any of them the ability to see any server or its internal state after the lockout procedure. Specifically, the server operators submit an agreed-upon program to the servers, which evaluate the program and publish the result to the server operators.

3 Secure Computation of an LLP Algorithm

In this section, we describe the LLP problem considered in the work of Comarela et al. [18], and then we describe our process to transform this into a secure multi-party computation protocol with data-oblivious control flow.

3.1 The LLP Problem

Learning from Label Proportions (LLP) is a weakly supervised machine learning problem in which the data is divided into groups or *bags* [57]. For each bag, only the proportion of labels is known, i.e., no individual label is available. The goal of LLP is to train a model that can correctly predict a label given a feature vector as input.

To be more specific, the machine learning task associated with this work can be characterized as follows. The input is composed of a set of pairs $D = \{(\mathbf{x}_i, b_i)\}_{i=1}^N$ and a vector \mathbf{p} , where: each \mathbf{x}_i is a feature vector, representing the web domains visited by an individual in a given period; b_i denotes the bag, or group, that the individual belongs to (in our work b_i is the individual's state of residence); and \mathbf{p} is a vector of proportions, indicating the proportion of individuals with preference for (without loss of generality) the Democratic Party candidate in each state. From this characterization, it is implied that each individual has a label: label 1 for a Democratic Party preference; and label 0 for a Republican Party preference. However, the individuals' labels are unknown. Hence, the goal of the LLP problem is to learn a model that can predict the label of individuals from D and \mathbf{p} . For the interested reader, a more detailed presentation of LLP and its possible variants can be found in the work of Franco et al. [26].

There are many algorithms for training machine learning models in LLP setups. In this work, we follow the approach presented by Comarela et al. [18], mainly for two reasons: first, it has been successfully used in similar contexts; and second, it is suitable for implementation in the employed MPC framework (as discussed in Sect. 3.2). In summary, the algorithm has three steps. The first step is label initialization, which can be performed randomly or according to proportions in each bag (i.e., U.S. state). The remaining two steps employ an iterative procedure that is repeated until convergence. In the second step, we train a logistic regression model using the previously assigned labels. Finally, in the third step, we refine the logistic regression model obtained in the last iteration and use it to compute updated predictions. More specifically, the default threshold (0.5) used to transform the model's probabilities into labels is changed on a per-bag basis, so that the predicted proportion of 1's in each bag (in the training set) matches the expected proportion (from the problem's input). We repeat the second and third steps until the predictions in consecutive iterations have converged.

The algorithm is formally defined in Algorithm 1, where we have combined Algs. 1 and 2 from Comarela et al. [18] for ease of presentation. As input, the algorithm takes the following:

- a data matrix \mathbf{U} , where each row is a user, and each column is a feature,
- a region map R , where $R(u)$ denotes the region to which a user u belongs,
- a proportion map B , where $B(r)$ denotes the fraction of users with label 1 for a region r , and
- a learning algorithm \mathcal{A} , which we instantiate with a logistic regression model.

As output, the algorithm produces a trained model Θ . The algorithm uses two labeling functions: $L(u)$ denotes the true label of a user u , while $L'(u)$ denotes our current prediction of the user’s label, and $P(L(u) = 1 \mid \Theta)$ is the predicted probability that a user has label 1 according to the model Θ .

Algorithm 1: Learning Algorithm (from Algorithms 1 and 2 of [18])

Input: $\mathbf{U}, R, B, \mathcal{A}$
Output: Θ

- 1 $L' \leftarrow \{\}$
- 2 **foreach** row $u \in U$ **do**
- 3 $L'(u) = \begin{cases} 1, & \text{if } B(R(u)) \geq 0.5 \\ 0, & \text{otherwise} \end{cases}$
- 4 **repeat**
- 5 $L'' \leftarrow L'$
- 6 $\Theta \leftarrow \text{train } \mathcal{A} \text{ on } (\mathbf{U}, L')$
- 7 **foreach** region $r \in R$ **do**
- 8 $t_r \leftarrow \text{percentile } (1 - B(r)) \text{ of } P(L(u) = 1 \mid \Theta) \forall u \in r$
- 9 **foreach** row $u \in \mathbf{U}$ **do**
- 10 $L'(u) = \begin{cases} 1, & \text{if } P(L(u) = 1 \mid \Theta) \geq t_{R(u)} \\ 0, & \text{otherwise} \end{cases}$
- 11 **until** $L' \approx L''$;
- 12 **return** Θ

3.2 Implementation Under MPC

Here we present a method of realizing Algorithm 1 under MPC, which is, to our knowledge, the first instantiation of a model for the LLP problem under MPC. We use and build upon the CrypTen [40] library. To simplify the following presentation, we freely and implicitly convert between an arithmetic secret sharing and a boolean secret sharing (in both directions) as needed. In practice, this is done using arithmetic-to-boolean and boolean-to-arithmetic conversions under MPC (see e.g., [22, 49, 55]).

Initialization (Lines 1–3). Initialization can be done in plaintext, as both the state proportions and the users’ reported states are known in the clear. One could imagine secret sharing the users’ states as well. In that case, initialization would entail obviously selecting the correct state proportion and obviously comparing it with 0.5.

Model Training (Line 6). Training the underlying model is one of the two cruxes of the computation. We instantiate the underlying learning algorithm \mathcal{A} with a logistic regression model. Logistic regression is supported out-of-the-box by CrypTen, allowing us to use standard features of the library to train Θ at each iteration.

Computing Thresholds (Lines 7–8). The other crux of the computation is determining the threshold for each state, because it requires oblivious sorting. Oblivious sorting is a complex and expensive operation under MPC. More importantly, it is not natively supported by CrypTen. We implement a vectorized Bitonic sort [7] on top of CrypTen that uses the library’s efficient primitives for batched operations to minimize round complexity. We sort each state’s predictions separately, meaning we never sort more than several hundred values in a single call to the sorting algorithm. As a result, we prioritize ease of implementation over maximizing performance.

Given an oblivious sorting implementation, computing thresholds becomes rather straightforward. We compute $P(L(u) = 1 \mid \Theta)$ for each row $u \in \mathbf{U}$ and sort the results. Since the state proportions are known in the clear, we can compute the target index in plaintext and retrieve the (secret-shared) result at that index in the list. We remark that it would certainly be possible to compute thresholds even if the state proportions were not known in the clear, but in this work we assume the proportions are public information and do not need to be protected.

Computing the New Predictions (Lines 9–10). Updating the predictions with the newly computed model and thresholds is straightforward. For all rows $u \in \mathbf{U}$, we compare the secret-shared inference result $P(L(u) = 1 \mid \Theta)$ with the secret-shared threshold. We can perform the comparisons in a single batch using a vectorized oblivious comparison. We store the secret-shared prediction bit for each user in L' .

Checking Convergence (Line 11). The last step in the algorithm is to check for convergence. We have two lists of secret-shared predictions, L' and L'' , and we need to determine their similarity. We first compute an elementwise (vectorized) oblivious equality to obtain a vector of secret-shared bits. The value we wish to compute is the number of those bits that are 1’s, and we wish to compare it to some convergence threshold. We sum the vector to obtain a secret sharing of the number of 1’s, and we obviously compare it with the convergence threshold. We open the comparison result, and if the value is greater than the threshold, the opened value is 1, so we terminate the algorithm. If it is 0, we continue to another iteration. The only information revealed by the opening is whether the model has converged or not.

In summary, we implement Algorithm 1 step-by-step while ensuring that each step follows a data-oblivious control flow. We use CrypTen out-of-the-box where possible and implement new functionality where needed.

4 Data and Results

In this section, we describe our case study application of the MPC-for-LLP technique to analyze web browsing data of several thousand participants during the months before and after the November 2024 U.S. presidential election.

4.1 Panelist Data and Metadata

We recruit users through an advertisement on Amazon Mechanical Turk, with an emphasis on recruiting panelists who live in “swing states”: that is, states that are likely to have close outcomes in the 2024 U.S. presidential election. Throughout the study, which lasted from August 2024 to February 2025, we collected data from 7926 total unique users, with over 2000 users contributing data daily for the majority of the study.

We collect both web browsing behavior and basic demographic information.

- In terms of web browsing data, the client browser extension contains a list of 517 popular websites, which we obtained by taking the union of the Alexa Top 500 list and the Semrush Top 100 list [61], removing advertising websites. The Alexa Top 500 list is no longer published, so we use the same version as Comarela et al. [18]. The extension records the number of visits that the panelist makes to the website, as well as the number of visits that resulted from referrals from one of the nine most popular social media websites. For all website visits, the extension records only the website’s domain, not a complete URL. On a daily basis, the browser uploads the histogram vector of web visits to the web server. Each data upload consists of the three encrypted secret shares of a vector of 1034 elements, along with plaintext metadata described next.
- In terms of demographic metadata, we collect the panelist’s Mechanical Turk ID, their self-reported state and ZIP code, the number of total visits to websites in the list, and the number of total referrals to websites in the list.

The Mechanical Turk ID is used to compensate the panelist for their participation. In each data upload, this ID is encrypted using message-locked encryption [9] that is only accessible to a member of our research team who processes the payments; this team member did not have access to any of the MPC compute servers. The MTurk ID is never provided as input to the machine learning algorithms discussed in Sect. 3, although the team member responsible for payment processing used the MTurk IDs to attach an anonymous ID to each upload for the purpose of anonymously aggregating data from the same user over several days. We discuss anonymous payments in Sect. 6.2 as a potential avenue for minimizing the trust required in the payment process.

The remaining metadata is used within the machine learning operation: the self-reported state is used to partition the data into bags for each state, and the total visits are used in order to convert counts into fractions. In principle, these operations could be performed under MPC as well (albeit with increased runtime). The state and ZIP code are also used within the data integrity check, described in Sect. 4.3.

4.2 Empirical Results

We use the panelist data together with the actual vote tally in the U.S. presidential election on November 5, 2024 to estimate political preferences at different points in time. As described in Sect. 3, our machine learning algorithm trains on the ground truth election results and corresponding web visits on November 5 and the surrounding days in order to produce model weights for the influence of each web domain on political preferences and state-by-state “cutoff” thresholds to assess candidate preferences. From this model, we use the web visit data to infer aggregate political preferences on other days, both backwards and forwards in time.

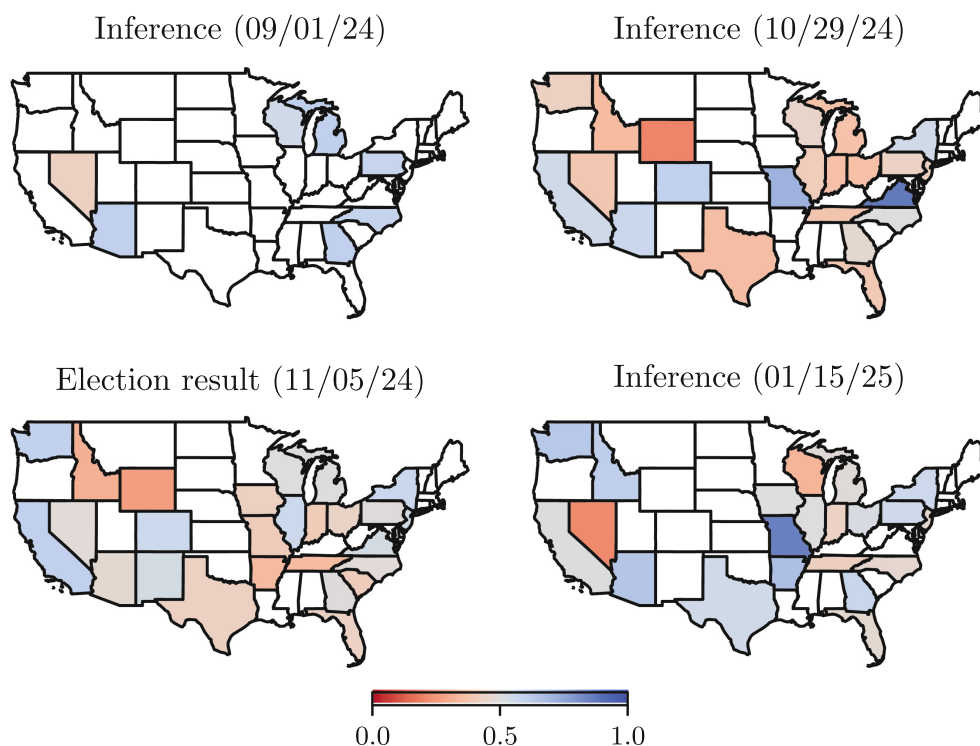


Fig. 2. State-level inference results at three points in time. One point is well before the election (top left), another is the week before the election (top right), and the last is after the election (bottom right). For comparison purposes, we also show the ground-truth election results (bottom left).

Findings. An example of our state-level inference results is shown in Fig. 2. Here we aggregate up the individual level candidate support scores for a set of states with ample respondents at each of three weekly periods: early in the campaign, election day, and the presidential inauguration. We have colored each state according to the degree of support for each party’s candidate, with red indicating more Republican and blue more Democratic.

While a detailed analysis of the results from a political science perspective is outside the scope of this work, we observe that our data show a general swing toward the Republican candidate before the election and a swing in the opposite direction after the election. Indeed, our indicators suggest a substantial movement towards the Republican candidate over the course of the campaign, with strong support in many of the swing states by the week of the election. In addition, our results suggest that the so-called honeymoon period of new presidents—where they typically enjoy high approval for the first few months of their presidency—was either totally missed or extremely short-lived in many states for the new president. The results suggest a number of avenues of future research for political science, which typically relies on polls that are limited in number and locales in these periods.

Performance. This case study also serves to demonstrate the viability of our methods for performing meaningful data analysis privately. An end-to-end training run including data for the week before and the week after the election takes approximately 70 min. Additionally, running inference for a single day’s data takes between 1 and 5 min, depending on the number of users in the sample on that day. All analysis took place on moderately powerful `m5n.2xlarge` AWS instances with 8 vCPUs and 32GB of memory. We emphasize that we did not seek to maximize performance, but only to make all computation run in a reasonable time frame of several hours. If we were to process, say, an order of magnitude more data, the present performance may become impractical, and more aggressive tuning and optimization would be necessary.

4.3 Data Integrity

For our inference of aggregate political preference to be accurate, it is critical that we correctly partition panelists into the bags for each state. For this reason, we ask panelists to self-report their state. However, we found that this self-reporting is not always accurate; for instance, our first Mechanical Turk advertisement was only open to people who live in swing states, which may have created an incentive for people to misrepresent their location. Although subsequent advertisements rewarded participants equally regardless of their reported state, the potential for dishonest reporting remains a concern. Future deployments should be careful to avoid incentive structures that encourage dishonest reporting. To examine the integrity of self-reported location information, we use geolocation data that allows us to draw conclusions about the validity of the sample.

Privacy Considerations. We exercised care to balance the data integrity goal with our overall privacy commitment. We design a validation process with the overarching objective of never having long-term storage of IP addresses or of location data beyond the ZIP code level. When the browser extension connects to the webserver to upload data, we use the connecting IP address and the IP locating service `ipgeolocation.io` to infer the panelist’s state and ZIP code and compare to the self-reported data. By performing this check at the time of upload, we do not need to store IP addresses.

Findings. Our primary takeaway is that while there are significant reasons to be concerned about the honesty of the users, the concerns ultimately are unlikely to substantially harm the ability to extract meaningful patterns from the data.




Fig. 3. Fractions of users who are honest and dishonest about their state of residence when filtering by a minimum number of days that the user is in the sample. The left and right plots consider a user honest if they are in the correct state for more than 0% or 90% of their total days in the sample.

Approximately 98% of the uploads come from users located in the United States, indicating that the vast majority of web browsing data corresponds to our target demographic at the highest level. However, the state-level integrity is much lower. Only 15% of uploads come from users located in their self-reported state. Furthermore, 15% of uploads contain a self-reported ZIP code that is not a valid ZIP code for the self-reported state. Taken together, the state-level integrity and ZIP code validity data paint a picture of a noisy dataset.

However, perhaps surprisingly, the honest panelists provide far more valuable data. To justify the data sample, we show that not all panelists and uploads are equal in terms of our ability to learn from them.

First, Fig. 3 shows that as we select for panelists who were in the sample for an increasingly large threshold of days, the panelists become more honest. We consider two ways of categorizing panelists into honest and dishonest groups: either we consider any panelist who has *ever* had their inferred state match their reported state to be honest, or we consider panelists to be honest if their inferred state matched their reported state for more than 90% of their uploads. In both cases, we see that the sample contains a large number of short-lived dishonest panelists, while consistent panelists tend to be more honest.

More importantly, Fig. 4 shows that honest panelists contribute visit and referral histograms that are significantly more dense than their dishonest counterparts. Using the same method of splitting panelists into honest and dishonest groups, we see that in either case, honest panelists and dishonest panelists contribute similar numbers of total visits, while honest panelists far outpace dishonest panelists in their total referrals. This observation is important due to the

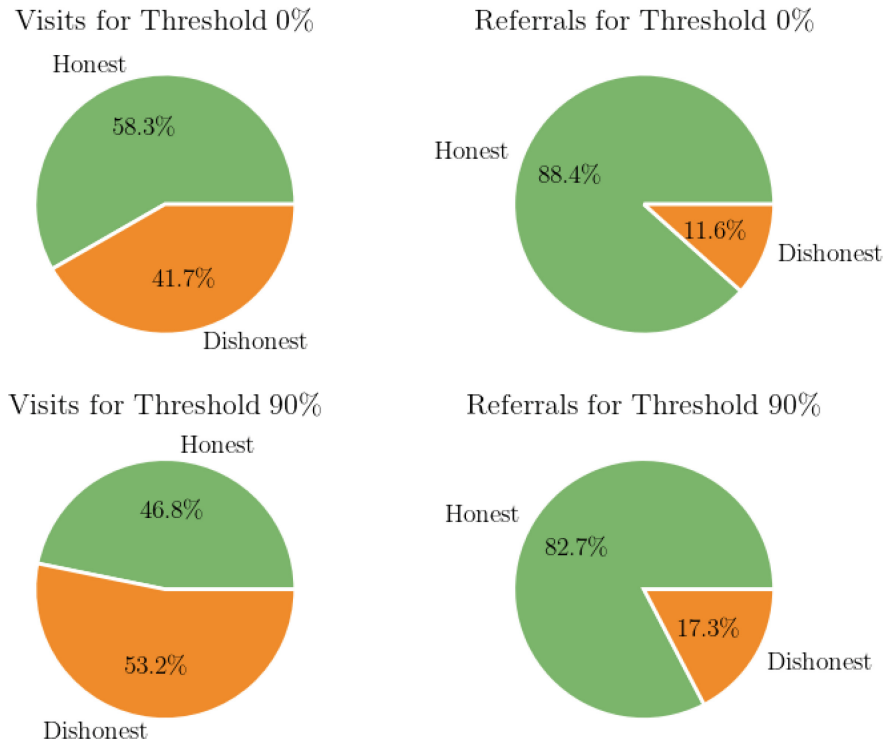


Fig. 4. The fractions of total visits and referrals reported by honest and dishonest users. The upper and lower plots consider a user honest if they are in the correct state for more than 0% or 90% of their total days in the sample.

finding by Comarela et al. [18] that referral data provides a far better signal for political preference than visits alone do.

We can draw two conclusions from the data and discussion above. First, detecting lying users and, if possible, obtaining some kind of proof or verification that a user is honest is of critical importance for ensuring data integrity, and this is an important direction for future work, as we discuss in Sect. 6. Second, while we should aim to maximize the integrity of the data, the fact that the heaviest users tend to be honest means our learning process is surprisingly robust to dishonest users.

5 Related Work

This work presents a deployment of privacy-preserving machine learning (PPML). In this section, we describe and compare with a few different categories of related work into PPML.

Algorithmic Designs. Building upon a long line of research into MPC in general, there have been significant advances in the past decade in the design of privacy-preserving machine learning algorithms. Some initial work, including SecureML [50], ABY³ [49], and SecureNN [66], built broad support for machine

learning operations under MPC. Next, several works contributed new protocols for privacy-preserving ML training and inference that collectively offer a variety of options for the number of parties, adversarial threshold, active vs. passive threat model, and more [15, 17, 19, 20, 37, 41, 42, 55, 56, 67]. A more recent line of work has explored privacy-preserving algorithms for deep learning operators like transformers [3, 23, 30, 31, 45, 48, 53, 72]. Our work only requires a relatively simple machine learning algorithm based on gradient descent (see Algorithm 1), which can be built even from the earlier line of PPML algorithms.

Software and Hardware Frameworks. There are several software frameworks for privacy-preserving machine learning. ABY³ [49] provides protocols for many specific fixed-point and vector operations used in machine learning. EzPC [16, 43, 59] provides a compiler from high-level machine learning operators to low-level crypto protocols. MP-SPDZ [39] is an extensive repository that contains several machine learning-friendly MPC protocols on its own right, and it is also used within the cloud-native Carbyne Stack framework [8].

Additionally, there are several frameworks that take advantage of dedicated hardware. Some works leverage dedicated GPUs [27, 33, 35, 36, 64, 69] or FPGAs [25, 34, 54] to improve performance. Other works combine MPC with trusted execution enclaves in order to provide defense-in-depth [21, 46, 70].

We use CrypTen [40] in this work due to its ease of use in developing Algorithm 1. That said, there are a number of frameworks listed above that we could have used instead, several of which also support the three-party setting that we use in this work.

Deployments of MPC. To the best of our knowledge, this work is the first use of MPC in the domain of political science. More generally, it contributes toward the ongoing efforts to increase adoption of MPC in scientific research settings to derive publicly useful findings from private data.

We provide several examples of prior work on MPC for public benefit. First, national statistics organizations have used MPC to conduct surveys and to aggregate data across agencies, as documented in a list compiled by the United Nations global working group on big data [65]. Second, policymakers have used MPC to gather statistical evidence for data-driven policymaking [2] in domains like education [12] and health services [32]. Third, healthcare organizations have used MPC to measure the aggregate spread of COVID-19 [5] and to conduct clinical research [60]. Finally, civic non-profit groups have used MPC to measure workplace culture [38] and empower survivors of sexual assault on college campuses [58].

6 Conclusion and Future Work

In this work, we have built and deployed a system for private analysis of political sentiment from web browsing data. We use the remainder of the paper to discuss some areas for possible improvement and future work based on lessons learned from the deployment.

6.1 Balancing Integrity Verification with Privacy

As with any data system, the quality of the data is of critical importance. Section 4.3 details our experience with measuring the honesty of users in reporting their residence, an important variable in our ability to make accurate state-level predictions. Our approach involved using the user’s IP address (which we will inevitably know by operating a standard webserver) to infer their current location. We could consider stronger methods of identity verification, but any strengthening would certainly lead to more invasive data collection, seemingly contradicting our core privacy goal.

In addition to the privacy element of the data integrity question, there is the fact that it is not immediately clear what information could be used to prove the location or residence of a user, even in a plaintext setting. Many companies and organizations use know-your-customer (KYC) systems for knowing the residence of their users. Perhaps one could design a private KYC service. We leave the problem of private residence verification as a challenging open problem.

6.2 Strengthening the Threat Model

The other major direction for future work is broadly to strengthen the threat model. To begin, the most immediate concern with the current system is the use of AWS as a single point of failure; this is an unfortunate but immediate consequence of our out-of-the-box use of CrypTen. While CrypTen provides a simple interface for running multiparty computation, future work could seek to build a similarly friendly interface that simultaneously can work across multiple cloud providers and multiple trust domains.

Another possible avenue towards a stronger threat model is to minimize the trust required from the user in the core machine learning computation. We can explore different system architectures and/or cryptographic primitives for the core computation. In the case of inference, the problem is comparatively easy, as each user can be treated independently, except for the aggregation of results. Given a model (either public or private), the user and a server could partake in a two-party computation to reveal only the prediction of the model on the user’s private input. Such a two-party computation would be lightweight enough to be run every day to obtain live-updated polling results. For training, the problem is more difficult, as individual users’ data can no longer be treated independently. One potential solution is to strengthen the MPC by incorporating multiple independent organizations to distribute the trust more widely. Alternatively, another option is to employ fully homomorphic encryption (FHE), involving limited client-side work to perform a threshold decryption to obtain the final results.

A final possible method for strengthening the threat model is the incorporation of anonymous payments. The payment method employed in this work requires knowledge of which Mechanical Turk IDs were active in the sample and for how many days. An anonymous payment mechanism would add an additional layer of protection for users beyond protecting their browsing data.

Acknowledgments. This research was supported by an Impact Program Grant from the Weidenbaum Center on the Economy, Government and Public Policy at Washington University; by the National Science Foundation under grants CNS-2209194, CNS-2312711, and CNS-2319369; by the DARPA SIEVE program under Agreement No. HR00112020021; and by a gift from Robert Bosch GmbH. Any opinions, findings, and conclusions or recommendations expressed in this material are those of the authors and do not necessarily reflect the views of the National Science Foundation or DARPA.

References

1. Abidin, A., Aly, A., Cleemput, S., Mustafa, M.A.: An MPC-based privacy-preserving protocol for a local electricity trading market. In: Foresti, S., Persiano, G. (eds.) CANS 2016. LNCS, vol. 10052, pp. 615–625. Springer, Cham (2016). https://doi.org/10.1007/978-3-319-48965-0_40
2. Advisory Committee on Data for Evidence Building. Year 2 report (2022). <https://www.bea.gov/sites/default/files/2022-10/acdeb-year-2-report.pdf>
3. Akimoto, Y., Fukuchi, K., Akimoto, Y., Sakuma, J.: Privformer: privacy-preserving transformer with MPC. In: 2023 IEEE European Symposium on Security and Privacy, pp. 392–410. IEEE Computer Society Press (2023)
4. Aono, Y., Hayashi, T., Phong, L.T., Wang, L.: Scalable and secure logistic regression via homomorphic encryption. In: Bertino, E., Sandhu, R., Pretschner, A. (eds.) Proceedings of the Sixth ACM on Conference on Data and Application Security and Privacy, CODASPY 2016, New Orleans, LA, USA, 9–11 March 2016, pp. 142–144. ACM (2016)
5. Apple and Google. Exposure notification privacy-preserving analytics (ENPA) white paper (2021). https://covid19-static.cdn-apple.com/-applications/covid19/current/static/contact-tracing/pdf/ENPA_White_Paper.pdf
6. Archer, D.W., et al.: From keys to databases – real-world applications of secure multi-party computation. Cryptology ePrint Archive, Report 2018/450 (2018)
7. Batcher, K.E.: Sorting networks and their applications. In: Proceedings of the April 30–May 2, 1968, Spring Joint Computer Conference, AFIPS 1968 (Spring), pp. 307–314. Association for Computing Machinery, New York (1968)
8. Becker, S., et al.: Carbyne Stack (2021). <https://github.com/carbynestack/carbynestack>
9. Bellare, M., Keelveedhi, S., Ristenpart, T.: Message-locked encryption and secure deduplication. In: Johansson, T., Nguyen, P.Q. (eds.) EUROCRYPT 2013. LNCS, vol. 7881, pp. 296–312. Springer, Heidelberg (2013). https://doi.org/10.1007/978-3-642-38348-9_18
10. Ben-Or, M., Goldwasser, S., Wigderson, A.: Completeness theorems for non-cryptographic fault-tolerant distributed computation (extended abstract). In: 20th ACM STOC, pp. 1–10. ACM Press (1988)
11. Bogdanov, D., Jõemets, M., Siim, S., Vaht, M.: How the Estonian tax and customs board evaluated a tax fraud detection system based on secure multi-party computation. In: Böhme, R., Okamoto, T. (eds.) FC 2015. LNCS, vol. 8975, pp. 227–234. Springer, Heidelberg (2015). https://doi.org/10.1007/978-3-662-47854-7_14
12. Bogdanov, D., Kamm, L., Kubo, B., Rebane, R., Sokk, V., Talviste, R.: Students and taxes: a privacy-preserving study using secure computation. PoPETs **2016**(3), 117–135 (2016)

13. Bogdanov, D., Laur, S., Willemson, J.: Sharemind: a framework for fast privacy-preserving computations. In: Jajodia, S., Lopez, J. (eds.) *ESORICS 2008*. LNCS, vol. 5283, pp. 192–206. Springer, Heidelberg (2008). https://doi.org/10.1007/978-3-540-88313-5_13
14. Bogetoft, P., et al.: Secure multiparty computation goes live. In: Dingledine, R., Golle, P. (eds.) *FC 2009*. LNCS, vol. 5628, pp. 325–343. Springer, Heidelberg (2009). https://doi.org/10.1007/978-3-642-03549-4_20
15. Byali, M., Chaudhari, H., Patra, A., Suresh, A.: FLASH: fast and robust framework for privacy-preserving machine learning. *PoPETs* **2020**(2), 459–480 (2020)
16. Chandran, N., Gupta, D., Rastogi, A., Sharma, R., Tripathi, S.: EzPC: programmable and efficient secure two-party computation for machine learning. In: *2019 IEEE European Symposium on Security and Privacy*, pp. 496–511. IEEE Computer Society Press (2019)
17. Chaudhari, H., Rachuri, R., Suresh, A.: Trident: efficient 4PC framework for privacy preserving machine learning. In: *NDSS 2020*. The Internet Society (2020)
18. Comarela, G., Durairajan, R., Barford, P., Christenson, D., Crovella, M.: Assessing candidate preference through web browsing history. In: *Proceedings of the 24th ACM SIGKDD International Conference on Knowledge Discovery & Data Mining*, pp. 158–167 (2018)
19. Dalskov, A.P.K., Escudero, D., Keller, M.: Secure evaluation of quantized neural networks. *PoPETs* **2020**(4), 355–375 (2020)
20. Dalskov, A.P.K., Escudero, D., Keller, M.: Fantastic four: honest-majority four-party secure computation with malicious security. In: Bailey, M., Greenstadt, R. (eds.) *USENIX Security 2021*, pp. 2183–2200. USENIX Association (2021)
21. de Laage, R., Yuhala, P., Wicht, F.-X., Felber, P., Cachin, C., Schiavoni, V.: Practical secure aggregation by combining cryptography and trusted execution environments. *CoRR*, abs/2504.08325 (2025)
22. Demmler, D., Schneider, T., Zohner, M.: ABY - a framework for efficient mixed-protocol secure two-party computation. In: *NDSS 2015*. The Internet Society (2015)
23. Dong, Y., et al.: Puma: Secure inference of llama-7b in five minutes (2023)
24. El Emam, K., et al.: A secure protocol for protecting the identity of providers when disclosing data for disease surveillance. *J. Am. Med. Inf. Assoc. JAMIA* **18**(3), 212–217 (2011)
25. Fang, X., Ioannidis, S., Leeser, M.: Secure function evaluation using an FPGA overlay architecture. In: *FPGA*, pp. 257–266. ACM (2017)
26. Franco, G., Crovella, M., Comarela, G.: Dependence and model selection in LLP: the problem of variants. In: *Proceedings of the 29th ACM SIGKDD Conference on Knowledge Discovery and Data Mining*, pp. 470–481 (2023)
27. Frederiksen, T.K., Jakobsen, T.P., Nielsen, J.B.: Faster maliciously secure two-party computation using the GPU. In: Abdalla, M., De Prisco, R. (eds.) *SCN 2014*. LNCS, vol. 8642, pp. 358–379. Springer, Cham (2014). https://doi.org/10.1007/978-3-319-10879-7_21
28. Gascón, A., et al.: Privacy-preserving distributed linear regression on high-dimensional data. *PoPETs* **2017**(4), 345–364 (2017)
29. Goldreich, O., Micali, S., Wigderson, A.: How to play any mental game or A completeness theorem for protocols with honest majority. In: Aho, A. (ed.) *19th ACM STOC*, pp. 218–229. ACM Press (1987)
30. Gupta, K., et al.: SIGMA: secure GPT inference with function secret sharing. *PoPETs* **2024**(4), 61–79 (2024)

31. Hao, M., Li, H., Chen, H., Xing, P., Xu, G., Zhang, T.: Private inference on transformers. In: NeurIPS, Iron (2022)
32. Hart, N.R., Archer, D.W., Dalton, E.: Privacy-preserved data sharing for evidence-based policy decisions: a demonstration project using human services administrative records for evidence-building activities (2019). <https://bipartisanpolicy.org/download/?file=/wp-content/uploads/2019/06/Privacy-Preserved-Data-Sharing-for-Evidence-Based-Policy-Decisions.pdf>
33. Harth-Kitzerow, C., Wang, Y., Rajat, R., Carle, G., Annavaram, M.: PIGEON: a high throughput framework for private inference of neural networks using secure multiparty computation. Cryptology ePrint Archive, Paper 2024/1371 (2024)
34. Huang, K., GÜngör, M., Fang, X., Ioannidis, S., Leeser, M.: Garbled circuits in the cloud using FPGA enabled nodes. In: HPEC, pp. 1–6. IEEE (2019)
35. Husted, N., Myers, S.A., Shelat, A., Grubbs, P.: GPU and CPU parallelization of honest-but-curious secure two-party computation. In: ACSAC, pp. 169–178. ACM (2013)
36. Jawalkar, N., Gupta, K., Basu, A., Chandran, N., Gupta, D., Sharma, R.: Orca: FSS-based secure training and inference with GPUs. In: 2024 IEEE Symposium on Security and Privacy, pp. 597–616. IEEE Computer Society Press (2024)
37. Juvekar, C., Vaikuntanathan, V., Chandrakasan, A.: GAZELLE: a low latency framework for secure neural network inference. In: Enck, W., Felt, A.P. (eds.) USENIX Security 2018, pp. 1651–1669. USENIX Association (2018)
38. Kaptchuk, G., Benoit-Bryan, J., Albab, K.D., Locks, M., Varia, M.: The good, the bad, and the ugly—lessons from an MPC for social good deployment. In: Real World Crypto Symposium (2024)
39. Keller, M.: MP-SPDZ: a versatile framework for multi-party computation. In: Ligatti, J., Ou, X., Katz, J., Vigna, G. (eds.) ACM CCS 2020, pp. 1575–1590. ACM Press (2020)
40. Knott, B., Venkataraman, S., Hannun, A.Y., Sengupta, S., Ibrahim, M., van der Maaten, L.J.P.: Crypten: secure multi-party computation meets machine learning. In: Proceedings of the NeurIPS Workshop on Privacy-Preserving Machine Learning (2020)
41. Koti, N., Pancholi, M., Patra, A., Suresh, A.: SWIFT: super-fast and robust privacy-preserving machine learning. In: Bailey, M., Greenstadt, R. (eds.) USENIX Security 2021, pp. 2651–2668. USENIX Association (2021)
42. Koti, N., Patra, A., Rachuri, R., Suresh, A.: Tetrad: actively secure 4PC for secure training and inference. In: NDSS 2022. The Internet Society (2022)
43. Kumar, N., Rathee, M., Chandran, N., Gupta, D., Rastogi, A., Sharma, R.: CryptFlow: secure TensorFlow inference. In: 2020 IEEE Symposium on Security and Privacy, pp. 336–353. IEEE Computer Society Press (2020)
44. Lapets, A., et al.: Accessible privacy-preserving web-based data analysis for assessing and addressing economic inequalities. In: COMPASS, pp. 48:1–48:5. ACM (2018)
45. Li, D., Wang, H., Shao, R., Guo, H., Xing, E., Zhang, H.: MPCFormer: fast, performant and private transformer inference with MPC. In: The Eleventh International Conference on Learning Representations (2023)
46. Li, X., Zhao, B., Yang, G., Xiang, T., Weng, J., Deng, R.H.: A survey of secure computation using trusted execution environments. CoRR, abs/2302.12150 (2023)
47. Liu, C., Wang, X.S., Nayak, K., Huang, Y., Shi, E.: OblivM: a programming framework for secure computation. In: 2015 IEEE Symposium on Security and Privacy, pp. 359–376. IEEE Computer Society Press (2015)

48. Lu, W., et al.: BumbleBee: secure two-party inference framework for large transformers. In: 32nd Annual Network and Distributed System Security Symposium, NDSS 2025. The Internet Society, 23–28 February 2025
49. Mohassel, P., Rindal, P.: ABY³: a mixed protocol framework for machine learning. In: Lie, D., Mannan, M., Backes, M., Wang, X. (eds.) ACM CCS 2018, pp. 35–52. ACM Press (2018)
50. Mohassel, P., Zhang, Y.: SecureML: a system for scalable privacy-preserving machine learning. In: 2017 IEEE Symposium on Security and Privacy, pp. 19–38. IEEE Computer Society Press (2017)
51. Nayak, K., Wang, X.S., Ioannidis, S., Weinsberg, U., Taft, N., Shi, E.: GraphSC: parallel secure computation made easy. In: 2015 IEEE Symposium on Security and Privacy, pp. 377–394. IEEE Computer Society Press (2015)
52. Nikolaenko, V., Weinsberg, U., Ioannidis, S., Joye, M., Boneh, D., Taft, N.: Privacy-preserving ridge regression on hundreds of millions of records. In: 2013 IEEE Symposium on Security and Privacy, pp. 334–348. IEEE Computer Society Press (2013)
53. Pang, Q., Zhu, J., Möllering, H., Zheng, W., Schneider, T.: BOLT: privacy-preserving, accurate and efficient inference for transformers. In: 2024 IEEE Symposium on Security and Privacy, pp. 4753–4771. IEEE Computer Society Press (2024)
54. Patel, R., Wolfe, P.-F., Munafo, R., Varia, M., Herbordt, M.C.: Arithmetic and Boolean secret sharing MPC on FPGAs in the data center. In: HPEC, pp. 1–8. IEEE (2020)
55. Patra, A., Schneider, T., Suresh, A., Yalame, H.: ABY2.0: improved mixed-protocol secure two-party computation. In: Bailey, M., Greenstadt, R. (eds.) USENIX Security 2021, pp. 2165–2182. USENIX Association (2021)
56. Patra, A., Suresh, A.: BLAZE: blazing fast privacy-preserving machine learning. In: NDSS 2020. The Internet Society (2020)
57. Quadrianto, N., Smola, A.J., Caetano, T.S., Le, Q.V.: Estimating labels from label proportions. In: International Conference on Machine Learning, pp. 776–783 (2008)
58. Rajan, A., Qin, L., Archer, D.W., Boneh, D., Lepoint, T., Varia, M.: Callisto: a cryptographic approach to detecting serial perpetrators of sexual misconduct. In: COMPASS, pp. 49:1–49:4. ACM (2018)
59. Rathee, D., et al.: CryptFlow2: practical 2-party secure inference. In: Ligatti, J., Ou, X., Katz, J., Vigna, G. (eds.) ACM CCS 2020, pp. 325–342. ACM Press (2020)
60. Rogers, J., et al.: VaultDB: a real-world pilot of secure multi-party computation within a clinical research network. <https://arxiv.org/pdf/2203.00146>
61. Semrush. Top 100: The most visited websites in the us. <https://www.semrush.com/blog/most-visited-websites/>
62. Shamir, A.: How to share a secret. *Commun. Assoc. Comput. Mach.* **22**(11), 612–613 (1979)
63. Songhori, E.M., Hussain, S.U., Sadeghi, A.-R., Schneider, T., Koushanfar, F.: TinyGarble: highly compressed and scalable sequential garbled circuits. In: 2015 IEEE Symposium on Security and Privacy, pp. 411–428. IEEE Computer Society Press (2015)
64. Tan, S., Knott, B., Tian, Y., Wu, D.J.: CryptGPU: fast privacy-preserving machine learning on the GPU. In: 2021 IEEE Symposium on Security and Privacy, pp. 1021–1038. IEEE Computer Society Press (2021)
65. UN Global Working Group Task Team on Privacy Preserving Techniques. Case study repository (2024). <https://unstats.un.org/wiki/spaces/UGTTOPPT/pages/150012018/Case+study+repository>

66. Wagh, S., Gupta, D., Chandran, N.: SecureNN: 3-party secure computation for neural network training. *PoPETs* **2019**(3), 26–49 (2019)
67. Wagh, S., Tople, S., Benhamouda, F., Kushilevitz, E., Mittal, P., Rabin, T.: Falcon: honest-majority maliciously secure framework for private deep learning. *PoPETs* **2021**(1), 188–208 (2021)
68. Walsh, J.M., Varia, M., Cohen, A., Sellars, A., Bestavros, A.: Multi-regulation computing: examining the legal and policy questions that arise from secure multiparty computation. In: *CSLAW*, pp. 53–65. ACM (2022)
69. Watson, J.-L., Wagh, S., Popa, R.A.: Piranha: a GPU platform for secure computation. In: Butler, K.R.B., Thomas, K. (eds.) *USENIX Security 2022*, pp. 827–844. USENIX Association (2022)
70. Wu, P., Ning, J., Shen, J., Wang, H., Chang, E.-C.: Hybrid trust multi-party computation with trusted execution environment. In: *NDSS 2022*. The Internet Society (2022)
71. Yao, A.C.-C.: Protocols for secure computations (extended abstract). In: *23rd FOCS*, pp. 160–164. IEEE Computer Society Press (1982)
72. Zeng, C., He, D., Feng, Q., Yang, X., Luo, Q.: SecureGPT: a framework for multi-party privacy-preserving transformer inference in GPT. *IEEE Trans. Inf. Forensics Secur.* **19**, 9480–9493 (2024)