

How Much User Context Do We Need? Privacy by Design in Mental Health NLP Applications

Ramit Sawhney^{1,4*}, Atula Neerkaje^{1,5*}, Ivan Habernal³, Lucie Flek^{1,2}

¹ Conversational AI and Social Analytics (CAISA) Lab, University of Marburg, Germany

² Conversational AI and Social Analytics (CAISA) Lab, University of Bonn, Germany

³ Trustworthy Human Language Technologies,
Department of Computer Science, Technical University of Darmstadt, Germany

⁴ Mohamed bin Zayed University of Artificial Intelligence, UAE

⁵ Manipal Institute of Technology, India

ramit.sawhney@mbzuai.ac.ae, lucie.flek@bit.uni-bonn.de

Abstract

Clinical NLP tasks such as mental health assessment from text, must take social constraints into account - the performance maximization must be constrained by the utmost importance of guaranteeing privacy of user data. Consumer protection regulations, such as GDPR, generally handle privacy by restricting data availability, such as requiring to limit user data to 'what is necessary' for a given purpose. In this work, we reason that providing stricter formal privacy guarantees, while increasing the volume of user data in the model, in most cases increases benefit for all parties involved, especially for the user. We demonstrate our arguments on two existing suicide risk assessment datasets of Twitter and Reddit posts. We present the first analysis juxtaposing user history length and differential privacy budgets and elaborate how modeling additional user context enables utility preservation while maintaining acceptable user privacy guarantees.

Introduction

A growing body of work shows that NLP models can effectively support professional human counselors in identifying mental health risk markers in online user behavior, for example to aid suicide risk assessment (McCarthy 2010; De Choudhury et al. 2016; Reger, Stanley, and Joiner 2020; Shing et al. 2018). An early intervention by a counselor can be crucial, as 80% of patients at suicidal risk do not undergo medical treatment. On the contrary, people turn to online platforms, with 8 out of 10 people disclosing their suicidal plans (Golden, Weiland, and Peterson 2009; Robinson et al. 2016). It is therefore desirable to train analytical models for these platforms, however, without unnecessarily exposing personal details of the patients in the training data.

State-of-the-art models for this task typically employ large contextual models (Matero et al. 2019; Losada, Crestani, and Parapar 2019; Zirikly et al. 2019). Leveraging additional data, such as user's history on social media, augments the predictive power (Zirikly et al. 2019; Sawhney et al. 2021). However, as every machine learning model, these can be prone to learning undesired data artifacts, for example users mentioning certain locations, per-

sons, or other rare pieces of information can be systematically misclassified as suicidal.

Moreover, privacy risks in such NLP models might lead to severe individual consequences. Not only can adversaries exploit sensitive artifacts from the models, but in extreme cases they can target vulnerable users (Hsin, Torous, and Roberts 2016).

Although consumer data laws address protection of vulnerable users to some extent, e.g., by limiting the amount of necessary personal data for processing, data minimization does not ensure increased privacy protection. In contrary, this can be even counterproductive, as reducing user data may lead to less robust classifiers and more artifacts learned. At the same time, advanced approaches such as *differentially private learning* (Dwork and Roth 2013) can be taken to abstain from using the privacy-breaching information in the learning process. However, these approaches typically lead to lower predictive power, making the classifiers less useful in practice (Alvim et al. 2011).

This raises an obvious empirical research question - **how much user data shall we use?** If data minimization does not increase privacy protection of users, can we include *more* data instead and protect its privacy by design? If so, which decisions do we need to consciously make in the model design? **What are the trade-offs between increasing formal user privacy preservation guarantees, maintaining predictive power, and minimizing users' training data?** Is there a way to find optimal thresholds for these?

In this paper, we present the first empirical analysis juxtaposing user history length and differential privacy budgets, demonstrating how modeling additional user context can compensate the performance loss while keeping acceptable privacy guarantees. Specifically, we demonstrate our approach on two suicide risk assessment datasets of Twitter and Reddit posts, as these tasks are vulnerable to user harm by privacy leaks (O'Loughlin et al. 2019; Rubanovich, Zisook, and Bloss 2022; Mikal, Hurst, and Conway 2016). We qualitatively inspect data points where enforcing differential privacy leads to performance changes, quantitatively estimate the privacy leakage by adversarial attacks, and examine the impact of different levels of class imbalance. Based on our experiments, we argue that

*Equal contribution

providing stricter formal privacy guarantees, while increasing the volume of user data in the model, in most cases increases the benefit for all parties involved, and that this effect shall be taken into account by future policy makers.

Background and Prior Work

A straight-forward approach to privacy protection in text is to identify sensitive user-revealing passages in a document and replace them with more general expressions (Hill et al. 2016; Alawad et al. 2020). However, removing personally identifiable information is typically not sufficient, as the summary statistics of the dataset provide means to infer individual’s membership with high probability (Jones et al. 2007; Sweeney 1997). Not only datasets, but also trained NLP models are vulnerable to attacks reconstructing the training data (Shokri et al. 2017). Privacy preservation efforts in NLP mostly follow adversarial approaches to obscure sensitive user characteristics at training time, such that the representations learned are invariant to these attributes (Li, Baldwin, and Cohn 2018; Friedrich et al. 2019; Coavoux, Narayan, and Cohen 2018). However, this approach offers only empirical privacy improvements, without any formal guarantees.

Differential Privacy (DP) has been adopted as a standard framework to provide a quantifiable privacy guarantee (Dwork and Roth 2013). Intuitively, a randomized algorithm is differentially private if the output distributions from two neighboring input databases, that are identical in all but a single training example, are indistinguishable (Fig. 1). The basic idea is that the output of the mechanism is “not much different” whether or not a single individual (example) is present in the dataset. The amount of “not much different” is controlled by the parameter ϵ , the privacy budget, and has a probabilistic interpretation of how much privacy of a single individual can be “lost” in the worst case.

In NLP, this paradigm is being mainly applied in the context of learning private word representations (Fernandes, Dras, and McIver 2019; Beigi et al. 2019; Coavoux, Narayan, and Cohen 2018) to prevent memorization in large language models (Carlini et al. 2019; Canfora et al. 2018) rather than directly in the downstream NLP classifiers, which we explore in this paper.

DP and performance drop. An increased privacy protection comes with a performance cost, typically leading to a decreased utility of the model (Alvim et al. 2011; Bassily, Smith, and Thakurta 2014; Shokri and Shmatikov 2015; Shokri et al. 2017), mainly due to memorization being prevented (van der Veen et al. 2018). While this effect has been described in machine learning research, its practical qualitative implications in the NLP field remain largely underexplored.

DP and disparate impact. Another challenge for applying privacy-preserving algorithms arises with data imbalance. Machine learning studies on other types of data (e.g. clinical records and images) show that the performance degradation is disparate, with minority subgroups of data suffering more utility loss (Shokri et al. 2017). Moreover, when stricter privacy guarantees are imposed, the utility gap widens disproportionately, i.e. the performance gap between

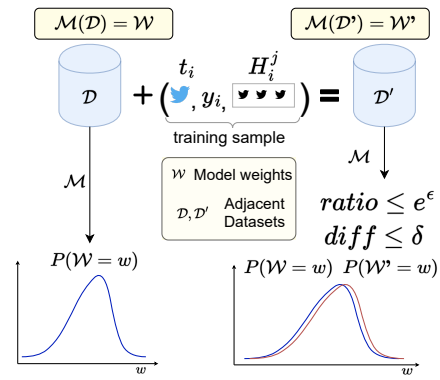


Figure 1: Essence of DP: With changes in a single training sample, the probability of obtaining a set of model weights remains approximately the same. Particularly, the ratio bounded by e^ϵ , and that difference bounded by δ .

the minority and majority subgroups is increasing (Bagdasaryan, Poursaeed, and Shmatikov 2019). In other words, the less represented groups which already have lower accuracy end up losing more utility. This disparate impact of DP can occur even where classes are only slightly imbalanced, and isn’t limited to strict privacy budgets (Farrand et al. 2020).

DP and the choice of privacy budget. The question of how to set the privacy budget (ϵ) has been present since the introduction of differential privacy and perceived as a “social question” (Dwork et al. 2006) with numerous mathematical answers proposed, such as utility-maximization optimizers (Geng and Viswanath 2014) and interval estimations (Naldi and D’Acquisto 2015). Voting mechanisms (Kohli and Laskowski 2018) gain popularity as they enable users to express their privacy budget; however, some studies find none to minimal links between privacy concerns reported by users and recorded behavior in a system with privacy warnings (Zimmerman et al. 2019). Some works argue for setting ϵ based on inferred user traits instead (Vu et al. 2017). In practice, values in the literature vary from as little as 0.01 to as much as 7, often with little to no justification. Recently, researchers examined DP settings of commercial companies (Tang et al. 2017), criticizing their generous ϵ choice. However, to date there is no clear consensus in the community on what strong or acceptable privacy budget values are. Generally, values under 1 can be considered strongly private, as a well-known DP method called Randomized response ($\epsilon = \ln(3)$) has been used for decades for highly sensitive questionnaires in social sciences (Warner 1965).

DP versus data minimization. Our motivation for contrasting the theoretical privacy preservation (as formalized by DP) with the data minimization approach originates from examining how the goal of personal data protection of the users is anchored in consumer data laws. For example the European Union’s General Data Protection Regulation (GDPR) requires limiting user data to *what is necessary concerning the purposes for which they are processed*,

encouraging data minimization rather than privacy protection by design. The operationalization of these laws in machine learning applications, including NLP, has been controversial (Tene and Polonetsky 2012; Biega et al. 2020), with practical implications of various design choices often remaining rather opaque to the engineers and policy makers alike. Several data minimization studies have been conducted in the domain of large recommender systems, noting that the overall performance often does not drop after removing random data points (Chow et al. 2013) or drastically reducing user history length (Wen et al. 2018). However, the incurred changes disparately impacted result quality for specific users or user groups (Biega et al. 2020). In this work, we conduct experiments to empirically examine such effects in the area of sensitive NLP tasks.

Methodology

We explore to which extent varying amounts of users’ historical data in combination with varying privacy budget impact performance of sensitive user-centric social NLP tasks, that is the suicide ideation and risk detection tasks.

Let S_i be the i -th training sample associated with user u_j from the set of M users. Each sample consists of a sequence of historical posts authored by the particular user u_j which we denote as $H_i^j = [h_1^i, h_2^i, \dots, h_L^i]$ which occur strictly before the time step where the prediction is made. The task is a classification problem, namely to predict label y_i .

We study two kinds of user-centric classification tasks. First, for *post-level* tasks, the suicidality label is associated with the most recent post to be assessed p_i . Here the user’s historical posts provide an additional context. The training sample S_i is hence the tuple (p_i, y_i, H_i^j) . Second, for *user-level* tasks, the suicidality label is associated with the user u_j , with the sample S_i representing the tuple (y_i, H_i^j) .

Tasks and Data Studied

We identify suicide ideation and risk detection on social media as prediction tasks which benefit from additional contextual data (Zirikly et al. 2019; Sawhney et al. 2021), with the need to protect private and sensitive data (Huckvale, Torous, and Larsen 2019; Rubanovich, Zisook, and Bloss 2022).

Suicide Ideation (SI) A post-level binary classification task, for which we use an existing Twitter dataset (Mishra et al. 2019) comprising Twitter timelines of 32,558 unique users, spanning over ten years of historical tweets from 2009 to 2019, summing up to 2.3M unlabeled tweets. These tweets were identified using a lexicon of 143 suicidal phrases. Out of these, 34,306 latest tweets were labeled as $y_i \in \{\textit{Suicidal Intent Absent}, \textit{Suicidal Intent Present}\}$. The dataset contains 3,984 suicidal tweets, indicating a high class imbalance of 9:1¹.

Suicide Risk – Reddit (SR) A user-level multi-class classification task, for which we use an existing dataset released by Gaur et al. (2019). The dataset consists of 9,127 Reddit posts by 500 users. Each user is associated with a suicide

risk severity label $y_i \in \{\textit{Indication}, \textit{Ideation}, \textit{Behaviour}, \textit{Attempt}\}$ in increasing order of suicide risk. The average number of posts made by a user is 18.25 ± 27.45 .²

User-Contextual NLP Modeling

We describe our base user-contextual model HistLSTM, which sequentially models user’s historical posts. For post-level tasks, features obtained by HistLSTM are combined with the post to be assessed before the classification step.

Encoding Posts We use the 768-dimensional representation of the [CLS] token obtained from BERT which tend to yield holistic representations of text, including posts on social media. We encode $\mathbf{P}'_i \in \mathbb{R}^{768}$ for a post to be assessed p_i as $\mathbf{P}'_i = \text{BERT}(p_i)$. Similarly, we encode each historical post h_k^i to yield the representation $\mathbf{E}_k^i \in \mathbb{R}^{768} = \text{BERT}(h_k^i)$.

Historical Context The sequential nature of historical posts makes LSTMs a natural method for contextual modeling. Each historical post embedding $\mathbf{E}_k^i \in \mathbb{R}^{768}$ is fed sequentially to the LSTM to obtain the final hidden state $\tilde{\mathbf{H}}_L^i \in \mathbb{R}^d = \text{LSTM}(\mathbf{E}^i)$. For post-level tasks, HistLSTM jointly learns from the language of the post to be assessed along with the sequentially modeled representation of the user’s timeline. For such tasks we apply the concatenation operation \oplus to \mathbf{P}'_i and $\tilde{\mathbf{H}}_L^i$ before feeding it to a dense layer with ReLU activation, followed by softmax.

$$\hat{\mathbf{y}}_i = \text{softmax}(\text{ReLU}(W_y(\mathbf{P}'_i \oplus \tilde{\mathbf{H}}_L^i) + b_y))$$

For user-level tasks, $\tilde{\mathbf{H}}_L^i$ is directly fed to dense layer:

$$\hat{\mathbf{y}}_i = \text{softmax}(\text{ReLU}(W_y(\tilde{\mathbf{H}}_L^i) + b_y))$$

where $\hat{\mathbf{y}}_i$ is the class probabilities vector and $\{W_y, b_y\}$ are trainable network parameters.

Differentially Private Optimization

Differential privacy has been adopted as a standard framework for privacy-preserving analysis. DP protects privacy of each single individual in the dataset by introducing random noise into a mechanism that queries the dataset. When a trained model is differentially private, the ϵ parameter limits the probability of an adversary to attack the model and reveal sensitive information from its training dataset. Formally, a mechanism \mathcal{M} with domain \mathcal{D} and range \mathcal{R} represented by $\mathcal{M} : \mathcal{D} \rightarrow \mathcal{R}$ satisfies (ϵ, δ) -differential privacy if for two adjacent datasets $d, d' \in \mathcal{D}$ that differ in one example and for any subset of outputs $\mathcal{S} \subseteq \mathcal{R}$ it holds that:

$$\Pr[\mathcal{M}(d) \in \mathcal{S}] \leq \exp(\epsilon) \cdot \Pr[\mathcal{M}(d') \in \mathcal{S}] + \delta \quad (1)$$

Lower ϵ implies better privacy. The term δ , which is typically ‘cryptographically small,’ is a relaxation of the ‘pure’ $(\epsilon, 0)$ -DP. It enables much better composition of several private algorithms in exchange for a small probability of privacy ‘failure’ (Abadi et al. 2016). In deep learning, the mechanism $\mathcal{M} : \mathcal{D} \rightarrow \mathcal{R}$ can be considered as a training

²We provide further insights on the individual class distributions in the appendix.

¹The imbalance is much greater in the real world

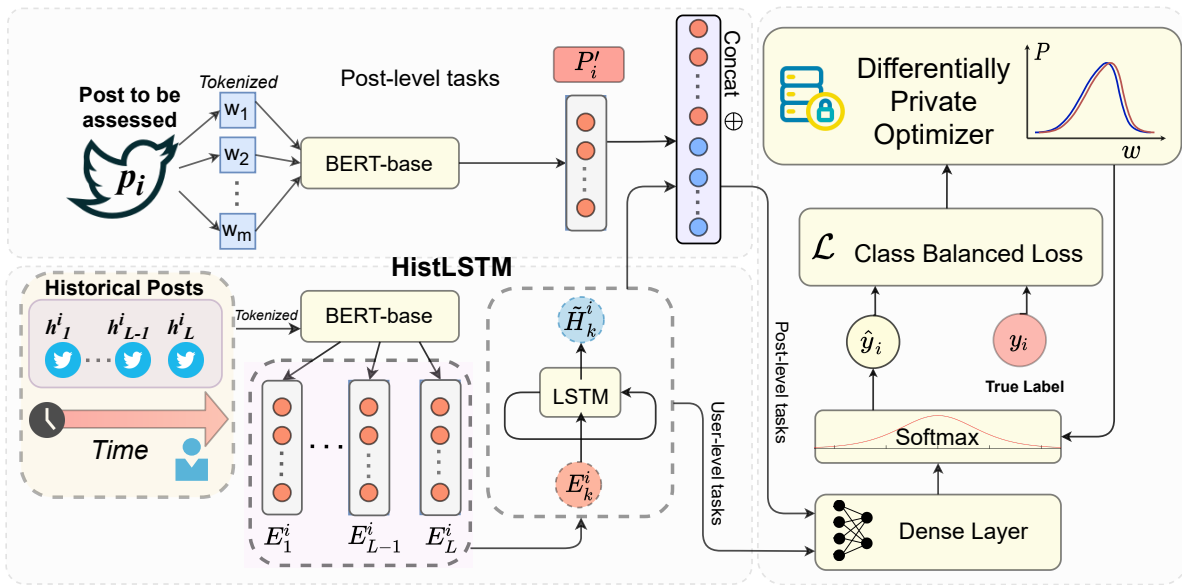


Figure 2: An overview of HistLSTM: We first encode all user posts. The historical post representations are sequentially modeled via an LSTM to obtain user historical context. For post-level tasks, it is concatenated with the post to be assessed and is used for classification. The model is optimized using a differentially private optimizer.

procedure \mathcal{M} on a dataset \mathcal{D} that returns a model in the space \mathcal{R} (Bagdasaryan, Poursaeed, and Shmatikov 2019).

Intuitively, Differential Privacy requires that the probability of learning any particular set of model parameters stays roughly the same if we change a single training example in the training set Figure 1. We implement differential privacy based on the widely adopted (Ha et al. 2019) DP-SGD algorithm introduced by Abadi et al. (2016).

The key aspects of differentially private SGD are (1) clipping the per-batch gradient where its norm exceeds the clipping bound C and (2) adding zero-centered Gaussian noise \mathcal{N} parametrized by σ to the aggregated per-sample gradients. Since the gradient computation is a function applied to the dataset and its output is privatized with DP, the resulting trained model will also be differentially-private. Intuitively, this means that by limiting the norm of the gradients and adding additional noise, we prevent the model from learning more than ϵ amount of information from every particular training sample, regardless of how distinct it is from the others. However, for DP-SGD, an important issue is quantifying the maximum privacy loss of each training sample i.e. computing the value of the privacy budget (ϵ).

We track the privacy budget (ϵ) using Rényi differential privacy (Mironov 2017), a relaxation of (ϵ, δ) -DP, which makes tracking the privacy budget easier. The moments accountant in Abadi et al. (2016) can be considered as a special case of Rényi differential privacy. For a given clipping bound C and Gaussian noise \mathcal{N} , the moments accountant “tracks” the privacy cost at each access to the training data, and accumulates this cost as the training progresses, eventually giving us the a final estimate of the privacy budget when training is stopped. We do not delve into those details here, but we point the reader to the appendix, and Abadi et al.

(2016) for extensive mathematical derivations of DP-SGD and the moments accountant.

Network Optimization and Classification

Often, in user-centric social NLP tasks, the posts of interest form only a small percentage of the data, leading to class imbalance. Thus, we apply Class-Balanced loss (Cui et al. 2019) to train HistLSTM, which introduces a weighting factor that is inversely proportional to the number of samples per class, yielding the loss \mathcal{L} as:

$$\mathcal{L}(\hat{y}_i, y_i) = \text{CB}_{focal}(\hat{y}_i, y_i; \beta, \gamma) \quad (2)$$

where CB_{focal} is class-balanced focal loss, y_i is the true label, and β and γ are hyperparameters.

We use a differentially private Adam optimizer, a variant of DP-SGD. This optimization procedure leverages gradient clipping and noise addition to calculate the model parameter update Δw in the loss landscape. The moments accountant simultaneously computes the privacy budget ϵ at the end of each epoch.

Evaluation Metrics

We choose the evaluation metrics following existing work on suicide ideation detection (Gaur et al. 2019). For the Suicide Ideation dataset, we use Recall on the minority class (*Suicidal Intent Present*) as it is important to not misclassify posts with suicidal intent. We also report Macro F1 score for overall performance. For the Suicide Risk dataset, we use Graded Recall, Graded Precision and FScore, which are calculated using altered definitions of Recall and Precision to account for ordinal nature of risk labels (Gaur et al. 2019).

Model	Architecture and Settings	Suicide Ideation			Suicide Risk		
		M-F1	R _s	P _s	F1	Gr. R	Gr. P
CurrentPostRF (Sawhney et al. 2018)	Language features: LIWC, POS tags, n-grams fed to a Random Forest classifier. Only uses most recent post.	0.54	0.51	0.47	0.47	0.49	0.45
CurrentPostLSTM (Coppersmith et al. 2018)	Single-post model which utilizes GloVe embeddings of the post fed to an LSTM. Only uses most recent post.	0.58	0.59	0.52	0.51	0.51	0.53
HistCNN (Gaur et al. 2019)	Concatenated post embeddings of all posts fed to a CNN	0.71*	0.59*	0.72 †	0.55*	0.51	0.69 †
HistDecay (Mathur et al. 2020)	Ensemble of BiLSTM+Attention with exponentially weighted embeddings of historical posts	0.73*	0.76 *	0.58*	0.56*	0.60*	0.57*
HistLSTM (Ours)	Our base model with no DP optimization	0.75 †	0.71*	0.65†	0.59 †	0.61 †	0.62†
HistLSTM + DP (Ours)	HistLSTM optimized with a strict privacy budget ($\epsilon = 0.74$)	0.68	0.61*	0.53*	0.56†	0.58*	0.57*

Table 1: Mean of results obtained over 15 independent runs. * and † indicate that the result is significantly ($p < 0.005$) better than CurrentPostLSTM (Non-contextual model), and HistDecay (SOTA) under Wilcoxon’s Signed Rank test respectively. Bold and italics indicates best and second best performance respectively.

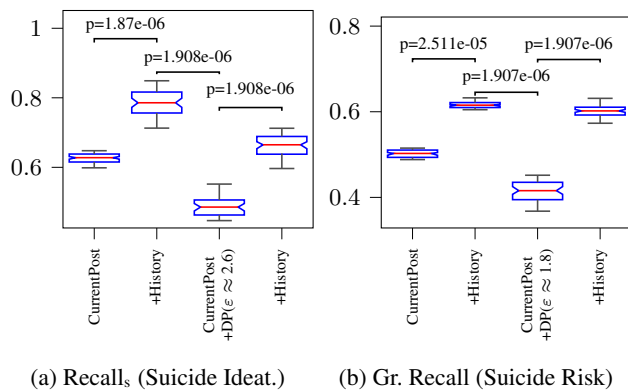


Figure 3: Confidence intervals for recall for the positive class over 10 independent runs for both tasks. (p) indicates the p -value under Wilcoxon’s Signed Rank test.

Experimental Setup

We split the twitter data using a temporal 70:10:20 (train:val:test) split, and the suicide risk dataset using a stratified 70:10:20 split. We use Macro F1/FScore on the validation set to select hyperparameters. Using grid search we explore: Hidden dimension H^d , Dropout d , β and γ , learning rate η , batch size N . We set δ as a value smaller than the inverse of the size of the dataset (Abadi et al. 2016). We manually set noise scale σ and clipping bound C . It is important to note that DP-SGD is not parameterized by the desired privacy budget ϵ , but rather by the noise scale σ and the clipping bound C , from which the final consumed privacy budget is computed. The general rule is that higher noise and lower clipping bounds yield tighter privacy budgets i.e. lower values of ϵ .

Results

What Is the Impact of User History Length?

In line with previous work (Khattari et al. 2015), we observe that adding historical context of a user timeline leads to significant ($p < 0.005$) improvement in performance for both tasks, as illustrated in Figure 3. With longer history, there is consistent increase in performance, as shown in Figure 4.

We additionally compare our model (HistLSTM) with other works that use single post only (CurrentPost), and those utilizing historical user context (Hist), in Table 1. As expected, user-contextual models provide significant improvements compared to single-post models (Dadvar et al. 2013). Among models with user history, temporally sequential models (HistDecay, HistLSTM) outperform bag-of-posts (HistCNN), likely due to their ability to capture the timeline build-up (Brådvik et al. 2008).

What Is the Impact of Privacy Budget Choice?

Incorporating differential privacy leads to performance drops, as shown in Figure 3. With less privacy (higher ϵ), the performance increases for a fixed window of historical volume (Figure 4). However, stricter private models ($\epsilon < 1$) with historical user context (HistLSTM + Strict DP) outperform non-private single-post methods, as shown in Table 1. For the SR dataset, strictly private models are on par with contextual models like HistCNN. Interestingly, models with a relaxed privacy budget perform on par with SOTA user-contextual methods, showing although there is a performance-privacy trade-off, there exist optimal balances where both aspects can be equitable.

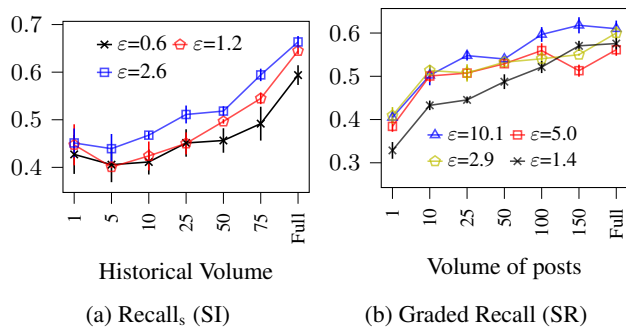


Figure 4: Performance (Recall) variations with increasing historical volume (number of historical tweets) and privacy budgets (ϵ) over 10 independent runs.

How Vulnerable to Privacy Attacks Are the Models?

To aid interpretability of the privacy budget and trade-offs involved, we present empirical results from conducting black-box membership inference attacks (Shokri et al. 2017) for various values of the privacy budget as shown in Table 2. We evaluate adversarial robustness using the privacy leakage metric (Jayaraman and Evans 2019), which is defined the difference between the true positive (TP) and false positive (FP) rate of the adversary. The key takeaways from our experiments are that data minimization alone is not sufficient to preserve privacy, reinforced by the observation that the non-private single-post model (*CurrentPost*) shows significant privacy leakage. The non-private HistLSTM performs best, at the cost of the highest privacy leakage. The incorporation of DP drastically reduces the adversarial efficacy with decreasing privacy budgets (ϵ), however at the cost of high performance degradation of the single-post model. Contrastingly, the addition of historical context to differentially private models increases performance while maintaining adversarial robustness. Specifically, even for strict privacy budgets ($\epsilon = 0.6$), HistLSTM shows significant utility, maintaining similar adversarial robustness as single-post models.

Can We Compensate Stricter Privacy Budget by Richer User History?

We observe adding historical context to the DP model leads to significant performance improvements ($p < 0.005$) with moderate DP budgets ($\epsilon \approx 2.6$, $\epsilon \approx 1.8$) (Figure 3). We note that DP models accounting for historical context outperform the non-contextual, non-private models in both tasks. Such models with user context and relaxed privacy budgets are able to match models with historical context (*HistDecay*), while models with strict privacy budgets ($\epsilon < 1$) outperform single-post methods (*Current*), as shown in Table 1.

We now contrast how performance varies over different user contexts (amount of historical volume) and privacy budgets in Figure 4. Specifically, we focus on Recall for the minority class (*Suicidal Intent Present*), and Graded Recall (SR dataset) as the metrics of interest (Figure 4). Also, for

Model	No DP		$\epsilon \approx 2.6$		$\epsilon \approx 0.6$	
	PL ↓	M. F1 ↑	PL ↓	M. F1 ↑	PL ↓	M. F1 ↑
CurrentPost	0.19	0.54	0.03	0.45	0.01	0.43
HistLSTM	0.38	0.75	0.05	0.69	0.02	0.62

Table 2: Model performance on the SI dataset with corresponding privacy leakage (PL) on conducting membership inference attacks for various values of privacy budget (ϵ).

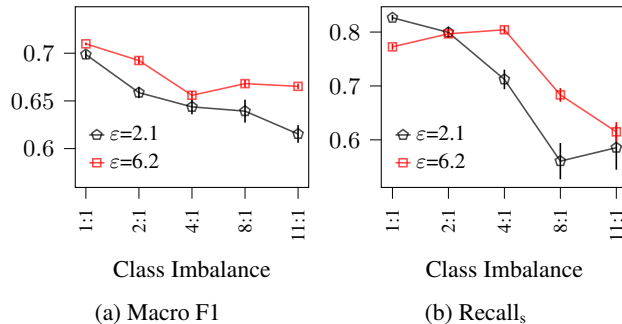


Figure 5: Changes in performance metrics with increasing class imbalance over 10 different runs (SI).

DP models, the performance for both tasks consistently increases with increasing historical volume. For more relaxed privacy budgets ($\epsilon > 2$), performance with full post history (over 50 tweets) is comparable to the non-private models. For lower privacy budgets the performance suffers significantly, yet surpassing non-private, single-post (*Current*) models (Table 1). The tradeoff between privacy and utility coupled with the historical data volume, is somewhat task-dependent. For the suicide ideation task, a more distant history (50 tweets) contributes to notable utility compensation (outperforming non-private counterparts with no historical context), even with strict privacy budgets, while the suicide risk model saturates earlier.

Probing the Impact of Class Imbalance

To study the effect of class imbalance on model performance and privacy, we train the model with varying class balances in Figure 5. For these experiments, we randomly sub-sample the full dataset before performing the split. We perform the random sub-sampling for each run of the experiments to minimize any sampling biases. We observe that for a given privacy budget (ϵ), the Macro F1 gradually decreases with increasing class imbalance. The recall of DP models drops much faster for the minority class when higher privacy guarantees are required, with drastic drops for the strict budgets ($\epsilon < 2.6$). Our observations align with existing work which suggest that privacy and fairness are fundamentally at odds under empirical risk minimization (Chang and Shokri 2021; Agarwal 2021; Tran, Dinh, and Fioretto 2021). Note that the performance degradation is disparate, with the underrepresented user groups (those at high suicide risk) suffering more utility loss. Interestingly, at the most extreme imbalance settings, we note slight performance improvements. Possibly

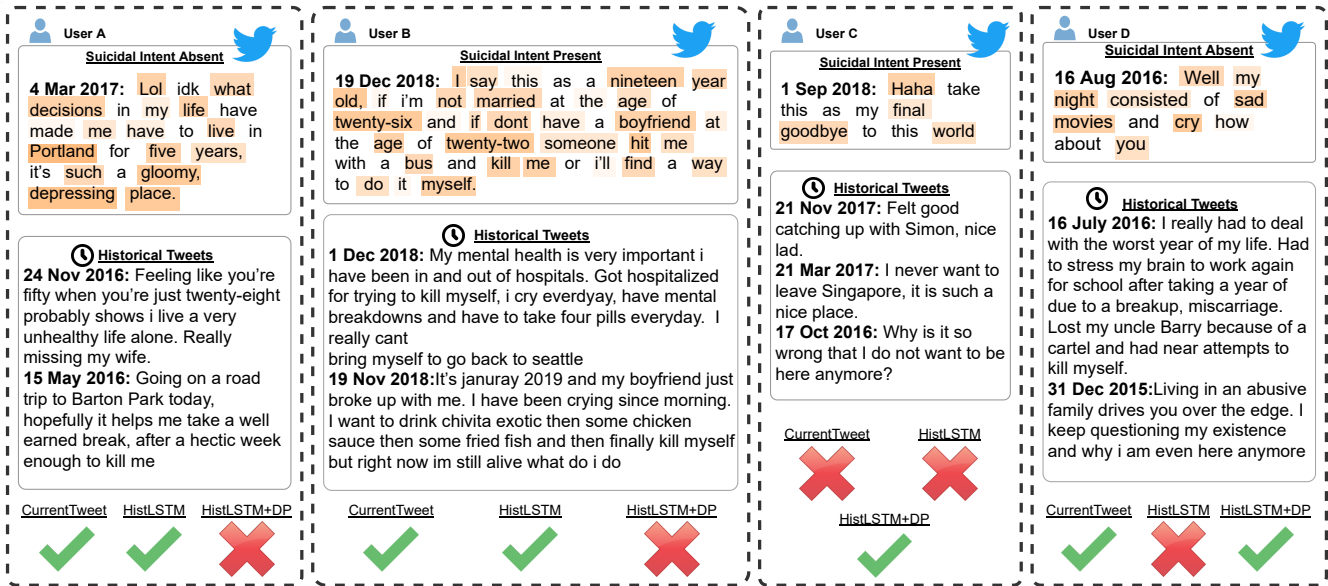


Figure 6: We study tweets of four users (SI dataset) and how a non-contextual model (CurrentPost) performs in comparison to a contextual model (HistLSTM, complete user history) and HistLSTM with strict ($\epsilon = 2.6$) privacy budget. We also visualize token self-attention of HistLSTM for descriptiveness (more intense means higher attention weight).

the tighter privacy guarantees degrade the utility on minority class so much that the model becomes more random, sometimes leading to an improvement (Farrand et al. 2020).

Qualitative Analysis

To add further insight beyond the challenging mathematical formulations of interpreting DP, we present a qualitative insight in Figure 6. Generally, we observe that augmenting models with DP leads to a failure in capturing relationships between the class labels and specific personal data related to named entities and numbers, which is sometimes desirable. For example, we see that tweets of User A and User B contain personal details such as locations and age. The non-private models likely leverage and memorize such details (Carlini et al. 2019), whereas the DP-based model misclassifies the tweet. For user C, models are thrown off by the positive words. Due to most of the historical tweets containing personal data, the tweet from 2016 may seem of higher importance to the DP variant, which correctly predicts the tweet to be suicidal.

Choice of Privacy Budget

In our experiments, we observe moderate privacy budgets ($0.5 < \epsilon < 4.5$) to notably decrease the sensitivity to names and digits (based on our subjective error analysis) while yielding acceptable performance when combined with full user timelines. Models with more benevolent (higher) ϵ choices tend to perform similarly to non-private settings.

Discussion

Trade-Offs: Privacy, Performance, Data Volume

Overall, settings where performance loss caused by high privacy guarantees is compensated by high data volumes, appear appealing and socially desirable. However, there might be cases in practice, where this choice is suboptimal, e.g. when highly imbalanced classes can result to socially and economically problematic disparate impact (Bagdasaryan, Poursaeed, and Shmatikov 2019). Furthermore, experiments of Shokri et al. (2017) also indicate disparate privacy vulnerability - smaller groups are more vulnerable to privacy disclosure under DP, which is in case of suicidal behavior a considerable design issue. Overall, an intervention in the form of stricter privacy budget (lower ϵ) can increase the risk of high outcome bias (error disparity), while a stricter data minimization may increase the risk of high input bias (selection bias or over-amplification) (Khani and Liang 2020).

Data minimization can be also consciously applied for purposes beyond privacy preservation, e.g. Vincent, Hecht, and Sen (2019) explore the “data strike” as a form of collective action of users against technology platforms, starving business-critical models of training data. The authors find the impact of data strikes rather limited. Moreover, the resulting data voids can be exploited by malicious users to bias the system (Golebiewski and Boyd 2018). Sometimes overall data minimization can be a prevalent goal for environmental reasons, as large models leave a major carbon footprint (Strubell, Ganesh, and McCallum 2019).

Conclusions

In this paper we have shown that leveraging *more* user-centric data enables competitive utility in prediction tasks

and, at the same time, protects individual's privacy, provided the training procedure uses differential privacy with a strict privacy budget. This setup increases the benefit for all parties involved, especially in sensitive tasks such as mental health assessment, where a strong privacy protection, hand in hand with low error rate, has been the most relevant factor for the patients themselves (Aledavood et al. 2017; Hsin, Torous, and Roberts 2016; Torous and Baker 2016). Strong privacy guarantees also enable sharing the research outcomes, such as deploying the trained models, in this sensitive domain (Lehman et al. 2021).

We present the first analysis juxtaposing user history length and differential privacy budgets. We demonstrate on two publicly available (Twitter and Reddit) datasets for suicide risk assessment how increasing the historical context used for each user helps compensate for the performance loss of DP models even with strict privacy budgets ($\epsilon \approx 0.6$), i.e., preserving high formal privacy guarantees.

We qualitatively inspect data points where enforcing differential privacy leads to performance drops and surprising performance improvements. We observe decreased importance of named entities and numerals in the model, such as the author's age and location, potentially leading to preferable, robust representations.

We show that increasing the imbalance of the predicted classes impacts the utility of DP models, given the disparate degradation in performance for minority classes, i.e., tweets with suicide intent. We also observe that a stricter choice of privacy budget results in more pronounced disparate degradation, i.e. steeper drop in recall for the minority class.

While we generally endorse applying private and data-minimizing model designs on user-centric NLP tasks, we point out that the disparate impact of such choice shall be reflected in design decisions.

Broader Impact & Ethical Considerations

Emphasizing the sensitive nature of this work, we acknowledge the trade-off between privacy and effectiveness. To avoid coercion and intrusive treatment, we work within the purview of acceptable privacy practices suggested by Chancellor et al. (2019). We paraphrase all examples shown in this work using the moderate disguise scheme (Bruckman 2002) to protect user privacy (Chancellor et al. 2019). We utilize publicly available data in a purely observational (Norval and Henderson 2017), and non-intrusive manner. For all tasks, we specifically only use existing datasets, and all user data is kept separately on protected servers linked to the raw text and network data only through anonymous IDs.

We acknowledge that suicidality is subjective, the interpretation of this analysis may vary across individuals on social media (Puschman 2017), and we do not know the true intentions of the user behind the post. Care should be taken so as to not to create stigma, and interventions must hence be carefully planned by consulting relevant stakeholders, such as clinicians, designers, and researchers (Chancellor et al. 2016). We acknowledge that suicide risk exists on a diverse spectrum (Bryan and Rudd 2006), and a binary distinction is a task simplification intended to alert the human in the loop about exceeding a possible intervention threshold. We note

that the studied data is limited to English-speaking Twitter, and also recognize that the data may be susceptible to other demographic, annotator, and medium-specific biases (Hovy and Spruit 2016). Although our work attempts to analyze aspects of users' nuanced and complex experiences, we acknowledge the limitations and potential misrepresentations that can occur when researchers analyze social media data, particularly data from a group to which the researchers do not explicitly belong.

We acknowledge that it is almost impossible to prevent abuse of released technology even when developed with good intentions (Hovy and Spruit 2016). Hence, we ensure that this analysis is shared only selectively and subject to IRB approval (Zimmer 2009) to avoid misuse such as Samaritan's Radar (Hsin, Torous, and Roberts 2016). Moreover, we aim to strive for an informed public, by addressing the dual-use threat with preemptive disclosure accompanying the code, in line with Solaiman et al. (2019). Our work does not make any diagnostic claims related to suicide. Our models and analysis should form part of a distributed human-in-the-loop system for finer interpretation of risk.

Acknowledgments

This work was supported by the German Federal Ministry of Education and Research (BMBF) as part of the Junior AI Scientists program under reference 01-S20060. The independent research group TrustHLT is supported by the Hessian Ministry of Higher Education, Research, Science and the Arts. This work was also supported by the National Research Center for Applied Cybersecurity ATHENE. We thank the anonymous reviewers for their valuable comments, which helped us improve the paper.

References

- Abadi, M.; Chu, A.; Goodfellow, I.; McMahan, H. B.; Mironov, I.; Talwar, K.; and Zhang, L. 2016. Deep learning with differential privacy. In *Proceedings of the 2016 ACM SIGSAC Conference on Computer and Communications Security*, 308–318.
- Agarwal, S. 2021. Trade-Offs between Fairness and Privacy in Machine Learning. In *IJCAI 2021 Workshop on AI for Social Good*.
- Alawad, M.; Yoon, H.-J.; Gao, S.; Mumphrey, B.; Wu, X.-C.; Durbin, E. B.; Jeong, J. C.; Hands, I.; Rust, D.; Coyle, L.; et al. 2020. Privacy-Preserving Deep Learning NLP Models for Cancer Registries. *IEEE Transactions on Emerging Topics in Computing*.
- Aledavood, T.; Hoyos, A. M. T.; Alakörkkö, T.; Kaski, K.; Saramäki, J.; Isometsä, E.; and Darst, R. K. 2017. Data collection for mental health studies through digital platforms: requirements and design of a prototype. *JMIR research protocols*, 6(6): e110.
- Alvim, M. S.; Andrés, M. E.; Chatzikokolakis, K.; Degano, P.; and Palamidessi, C. 2011. Differential privacy: on the trade-off between utility and information leakage. In *International Workshop on Formal Aspects in Security and Trust*, 39–54. Springer.

- Bagdasaryan, E.; Poursaeed, O.; and Shmatikov, V. 2019. Differential privacy has disparate impact on model accuracy. *Advances in Neural Information Processing Systems*, 32: 15479–15488.
- Bassily, R.; Smith, A.; and Thakurta, A. 2014. Private empirical risk minimization: Efficient algorithms and tight error bounds. In *2014 IEEE 55th Annual Symposium on Foundations of Computer Science*, 464–473. IEEE.
- Beigi, G.; Shu, K.; Guo, R.; Wang, S.; and Liu, H. 2019. Privacy Preserving Text Representation Learning. In *Proceedings of the 30th ACM Conference on Hypertext and Social Media*, HT '19, 275–276. New York, NY, USA: Association for Computing Machinery. ISBN 9781450368858.
- Biega, A. J.; Potash, P.; Daumé, H.; Diaz, F.; and Finck, M. 2020. Operationalizing the Legal Principle of Data Minimization for Personalization. In *Proceedings of the 43rd International ACM SIGIR Conference on Research and Development in Information Retrieval*, 399–408.
- Brådvik, L.; Mattisson, C.; Bogren, M.; and Nettelbladt, P. 2008. Long-term suicide risk of depression in the Lundby cohort 1947–1997—severity and gender. *Acta Psychiatrica Scandinavica*, 117(3): 185–191.
- Bruckman, A. 2002. Studying the amateur artist: A perspective on disguising data collected in human subjects research on the Internet. *Ethics and Information Technology*, 4(3).
- Bryan, C. J.; and Rudd, M. D. 2006. Advances in the assessment of suicide risk. *Journal of clinical psychology*, 62(2).
- Canfora, G.; Di Sorbo, A.; Emanuele, E.; Forootani, S.; and Visaggio, C. A. 2018. A nlp-based solution to prevent from privacy leaks in social network posts. In *Proceedings of the 13th International Conference on Availability, Reliability and Security*, 1–6.
- Carlini, N.; Liu, C.; Erlingsson, Ú.; Kos, J.; and Song, D. 2019. The secret sharer: Evaluating and testing unintended memorization in neural networks. In *28th {USENIX} Security Symposium ({USENIX} Security 19)*, 267–284.
- Chancellor, S.; Birnbaum, M. L.; Caine, E. D.; Silenzio, V. M. B.; and De Choudhury, M. 2019. A Taxonomy of Ethical Tensions in Inferring Mental Health States from Social Media. In *Proceedings of the Conference on Fairness, Accountability, and Transparency*, FAT* '19. New York, NY, USA: Association for Computing Machinery.
- Chancellor, S.; Lin, Z.; Goodman, E. L.; Zerwas, S.; and De Choudhury, M. 2016. Quantifying and Predicting Mental Illness Severity in Online Pro-Eating Disorder Communities. In *Proceedings of the 19th ACM Conference on Computer-Supported Cooperative Work & Social Computing*, CSCW '16, 1171–1184. New York, NY, USA: Association for Computing Machinery. ISBN 9781450335928.
- Chang, H.; and Shokri, R. 2021. On the privacy risks of algorithmic fairness. In *2021 IEEE European Symposium on Security and Privacy (EuroS&P)*, 292–303. IEEE.
- Chow, R.; Jin, H.; Knijnenburg, B.; and Saldamli, G. 2013. Differential data analysis for recommender systems. In *Proceedings of the 7th ACM conference on Recommender systems*, 323–326.
- Coavoux, M.; Narayan, S.; and Cohen, S. B. 2018. Privacy-preserving Neural Representations of Text. In *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing*, 1–10.
- Coppersmith, G.; Leary, R.; Crutchley, P.; and Fine, A. 2018. Natural language processing of social media as screening for suicide risk. *Biomedical informatics insights*, 10.
- Cui, Y.; Jia, M.; Lin, T.-Y.; Song, Y.; and Belongie, S. 2019. Class-balanced loss based on effective number of samples. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 9268–9277.
- Dadvar, M.; Trieschnigg, D.; Ordelman, R.; and de Jong, F. 2013. Improving cyberbullying detection with user context. In *European Conference on Information Retrieval*. Springer.
- De Choudhury, M.; Kiciman, E.; Dredze, M.; Coppersmith, G.; and Kumar, M. 2016. Discovering shifts to suicidal ideation from mental health content in social media. In *Proceedings of the 2016 CHI conference on human factors in computing systems*, 2098–2110.
- Dwork, C.; McSherry, F.; Nissim, K.; and Smith, A. 2006. Calibrating noise to sensitivity in private data analysis. In *Theory of cryptography conference*, 265–284. Springer.
- Dwork, C.; and Roth, A. 2013. The Algorithmic Foundations of Differential Privacy. *Foundations and Trends® in Theoretical Computer Science*, 9(3-4): 211–407.
- Farrand, T.; Mireshghallah, F.; Singh, S.; and Trask, A. 2020. Neither Private Nor Fair: Impact of Data Imbalance on Utility and Fairness in Differential Privacy. In *Proceedings of the 2020 Workshop on Privacy-Preserving Machine Learning in Practice*, 15–19.
- Fernandes, N.; Dras, M.; and McIver, A. 2019. Generalised Differential Privacy for Text Document Processing. In Nielson, F.; and Sands, D., eds., *Principles of Security and Trust*, 123–148. Cham: Springer International Publishing.
- Friedrich, M.; Köhn, A.; Wiedemann, G.; and Biemann, C. 2019. Adversarial Learning of Privacy-Preserving Text Representations for De-Identification of Medical Records. In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, 5829–5839. Florence, Italy: Association for Computational Linguistics.
- Gaur, M.; Alambo, A.; Sain, J. P.; Kursuncu, U.; Thirunarayan, K.; Kavuluru, R.; Sheth, A.; Welton, R.; and Pathak, J. 2019. Knowledge-aware assessment of severity of suicide risk for early intervention. In *The World Wide Web Conference*, 514–525.
- Geng, Q.; and Viswanath, P. 2014. The optimal mechanism in differential privacy. In *2014 IEEE international symposium on information theory*, 2371–2375. IEEE.
- Golden, R. N.; Weiland, C.; and Peterson, F. 2009. *The truth about illness and disease*. Infobase Publishing.
- Golebiewski, M.; and Boyd, D. 2018. Data voids: Where missing data can easily be exploited. *Data & Society*, 29.
- Ha, T.; Dang, T. K.; Dang, T. T.; Truong, T. A.; and Nguyen, M. T. 2019. Differential privacy in deep learning: an overview. In *2019 International Conference on Advanced Computing and Applications (ACOMP)*, 97–102. IEEE.

- Hill, S.; Zhou, Z.; Saul, L.; and Shacham, H. 2016. On the (in) effectiveness of mosaicing and blurring as tools for document redaction. *Proceedings on Privacy Enhancing Technologies*, 2016(4): 403–417.
- Hovy, D.; and Spruit, S. L. 2016. The social impact of natural language processing. In *Proceedings of the 54th Annual Meeting of the Association for Computational Linguistics (Volume 2: Short Papers)*, 591–598.
- Hsin, H.; Torous, J.; and Roberts, L. 2016. An Adjuvant Role for Mobile Health in Psychiatry. *JAMA Psychiatry*, 73(2).
- Huckvale, K.; Torous, J.; and Larsen, M. E. 2019. Assessment of the data sharing and privacy practices of smartphone apps for depression and smoking cessation. *JAMA network open*, 2(4): e192542–e192542.
- Jayaraman, B.; and Evans, D. 2019. Evaluating differentially private machine learning in practice. In *28th {USENIX} Security Symposium ({USENIX} Security 19)*.
- Jones, R.; Kumar, R.; Pang, B.; and Tomkins, A. 2007. "I know what you did last summer" query logs and user privacy. In *Proceedings of the sixteenth ACM conference on Conference on information and knowledge management*.
- Khani, F.; and Liang, P. 2020. Removing Spurious Features can Hurt Accuracy and Affect Groups Disproportionately. arXiv:2012.04104.
- Khattari, A.; Joshi, A.; Bhattacharyya, P.; and Carman, M. 2015. Your sentiment precedes you: Using an author's historical tweets to predict sarcasm. In *Proceedings of the 6th workshop on computational approaches to subjectivity, sentiment and social media analysis*, 25–30.
- Kohli, N.; and Laskowski, P. 2018. Epsilon voting: Mechanism design for parameter selection in differential privacy. In *2018 IEEE Symposium on Privacy-Aware Computing (PAC)*, 19–30. IEEE.
- Lehman, E.; Jain, S.; Pichotta, K.; Goldberg, Y.; and Wallace, B. C. 2021. Does BERT Pretrained on Clinical Notes Reveal Sensitive Data? arXiv preprint arXiv:2104.07762.
- Li, Y.; Baldwin, T.; and Cohn, T. 2018. Towards Robust and Privacy-preserving Text Representations. In *Proceedings of the 56th Annual Meeting of the Association for Computational Linguistics (Volume 2: Short Papers)*. Melbourne, Australia: Association for Computational Linguistics.
- Losada, D. E.; Crestani, F.; and Parapar, J. 2019. Overview of eRisk at CLEF 2019 Early Risk Prediction on the Internet (extended overview).
- Matero, M.; Idnani, A.; Son, Y.; Giorgi, S.; Vu, H.; Zamani, M.; Limbachiya, P.; Guntuku, S. C.; and Schwartz, H. A. 2019. Suicide risk assessment with multi-level dual-context language and bert. In *Proceedings of the Sixth Workshop on Computational Linguistics and Clinical Psychology*, 39–44.
- Mathur, P.; Sawhney, R.; Chopra, S.; Leekha, M.; and Shah, R. R. 2020. Utilizing Temporal Psycholinguistic Cues for Suicidal Intent Estimation. In *European Conference on Information Retrieval*, 265–271. Springer.
- McCarthy, M. J. 2010. Internet monitoring of suicide risk in the population. *Journal of affective disorders*, 122(3).
- Mikal, J.; Hurst, S.; and Conway, M. 2016. Ethical issues in using Twitter for population-level depression monitoring: a qualitative study. *BMC medical ethics*, 17(1): 1–11.
- Mironov, I. 2017. Rényi differential privacy. In *2017 IEEE 30th Computer Security Foundations Symposium (CSF)*, 263–275. IEEE.
- Mishra, R.; Prakhar Sinha, P.; Sawhney, R.; Mahata, D.; Mathur, P.; and Ratn Shah, R. 2019. SNAP-BATNET: Cascading Author Profiling and Social Network Graphs for Suicide Ideation Detection on Social Media. In *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Student Research Workshop*, 147–156. Minneapolis, Minnesota: Association for Computational Linguistics.
- Naldi, M.; and D'Acquisto, G. 2015. Differential privacy: an estimation theory-based method for choosing epsilon. arXiv preprint arXiv:1510.00917.
- Norval, C.; and Henderson, T. 2017. Contextual Consent: Ethical Mining of Social Media for Health Research. *CoRR*, abs/1701.07765.
- O'Loughlin, K.; Neary, M.; Adkins, E. C.; and Schueller, S. M. 2019. Reviewing the data security and privacy policies of mobile apps for depression. *Internet interventions*, 15: 110–115.
- Puschman, C. 2017. Bad Judgment, Bad Ethics? *Internet Research Ethics for the Social Age*, 95.
- Reger, M. A.; Stanley, I. H.; and Joiner, T. E. 2020. Suicide Mortality and Coronavirus Disease 2019—A Perfect Storm? *JAMA Psychiatry*.
- Robinson, J.; Cox, G.; Bailey, E.; Hetrick, S.; Rodrigues, M.; Fisher, S.; and Herrman, H. 2016. Social media and suicide prevention: a systematic review. *Early intervention in psychiatry*, 10(2): 103–121.
- Rubanovich, C. K.; Zisook, S.; and Bloss, C. S. 2022. Associations between privacy-related constructs and depression and suicide risk in health care professionals, trainees, and students. *Academic medicine*, 97(2): 239–246.
- Sawhney, R.; Joshi, H.; Flek, L.; and Shah, R. R. 2021. PHASE: Learning Emotional Phase-aware Representations for Suicide Ideation Detection on Social Media. In *Proc. EACL*, 2415–2428.
- Sawhney, R.; Manchanda, P.; Singh, R.; and Aggarwal, S. 2018. A computational approach to feature extraction for identification of suicidal ideation in tweets. In *Proceedings of ACL 2018, Student Research Workshop*, 91–98.
- Shing, H.-C.; Nair, S.; Zirikly, A.; Friedenber, M.; Daumé III, H.; and Resnik, P. 2018. Expert, Crowdsourced, and Machine Assessment of Suicide Risk via Online Postings. In *Proceedings of the Fifth Workshop on Computational Linguistics and Clinical Psychology: From Keyboard to Clinic*, 25–36. New Orleans, LA: Association for Computational Linguistics.
- Shokri, R.; and Shmatikov, V. 2015. Privacy-preserving deep learning. In *Proceedings of the 22nd ACM SIGSAC conference on computer and communications security*.

- Shokri, R.; Stronati, M.; Song, C.; and Shmatikov, V. 2017. Membership inference attacks against machine learning models. In *2017 IEEE Symposium on Security and Privacy (SP)*, 3–18. IEEE.
- Solaiman, I.; Brundage, M.; Clark, J.; Askell, A.; Herbert-Voss, A.; Wu, J.; Radford, A.; Krueger, G.; Kim, J. W.; Kreps, S.; et al. 2019. Release strategies and the social impacts of language models. *arXiv preprint arXiv:1908.09203*.
- Strubell, E.; Ganesh, A.; and McCallum, A. 2019. Energy and Policy Considerations for Deep Learning in NLP. In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, 3645–3650. Florence, Italy: Association for Computational Linguistics.
- Sweeney, L. 1997. Weaving technology and policy together to maintain confidentiality. *The Journal of Law, Medicine & Ethics*, 25(2-3): 98–110.
- Tang, J.; Korolova, A.; Bai, X.; Wang, X.; and Wang, X. 2017. Privacy loss in apple’s implementation of differential privacy on macos 10.12. *arXiv preprint arXiv:1709.02753*.
- Tene, O.; and Polonetsky, J. 2012. Big data for all: Privacy and user control in the age of analytics. *Nw. J. Tech. & Intell. Prop.*, 11: xxvii.
- Torous, J.; and Baker, J. T. 2016. Why psychiatry needs data science and data science needs psychiatry: connecting with technology. *JAMA psychiatry*, 73(1): 3–4.
- Tran, C.; Dinh, M.; and Fioretto, F. 2021. Differentially Private Empirical Risk Minimization under the Fairness Lens. *Advances in Neural Information Processing Systems*, 34.
- van der Veen, K. L.; Seggers, R.; Bloem, P.; and Patrini, G. 2018. Three tools for practical differential privacy. *arXiv preprint arXiv:1812.02890*.
- Vincent, N.; Hecht, B.; and Sen, S. 2019. Data Strikes’: Evaluating the Effectiveness of New Forms of Collective Action Against Technology Platforms. In *Proceedings of The Web Conference 2019*.
- Vu, X.-S.; Jiang, L.; Brändström, A.; and Elmroth, E. 2017. Personality-based knowledge extraction for privacy-preserving data analysis. In *Proceedings of the Knowledge Capture Conference*, 1–4.
- Warner, S. L. 1965. Randomized response: A survey technique for eliminating evasive answer bias. *Journal of the American Statistical Association*, 60(309): 63–69.
- Wen, H.; Yang, L.; Sobolev, M.; and Estrin, D. 2018. Exploring recommendations under user-controlled data filtering. In *Proceedings of the 12th ACM Conference on Recommender Systems*, 72–76.
- Zimmer, M. 2009. Web search studies: Multidisciplinary perspectives on web search engines. In *International handbook of internet research*, 507–521. Springer.
- Zimmerman, S.; Thorpe, A.; Fox, C.; and Kruschwitz, U. 2019. Investigating the Interplay Between Searchers’ Privacy Concerns and Their Search Behavior. In *Proceedings of the 42nd International ACM SIGIR Conference on Research and Development in Information Retrieval*, 953–956.
- Zirikly, A.; Resnik, P.; Uzuner, Ö.; and Hollingshead, K. 2019. CLPsych 2019 Shared Task: Predicting the Degree of Suicide Risk in Reddit Posts. In *Proceedings of the Sixth Workshop on Computational Linguistics and Clinical Psychology*, 24–33. Minneapolis, Minnesota: Association for Computational Linguistics.