

# Multimodal Privacy-preserving Mood Prediction from Mobile Data: A Preliminary Study

Terrance Liu<sup>1\*</sup>, Paul Pu Liang<sup>1\*</sup>, Michal Muszynski<sup>1</sup>, Ryo Ishii<sup>1</sup>,  
David Brent<sup>2</sup>, Randy P. Auerbach<sup>3</sup>, Nicholas B. Allen<sup>4</sup>, Louis-Philippe Morency<sup>1</sup>

<sup>1</sup>School of Computer Science, Carnegie Mellon University

<sup>2</sup>Department of Psychiatry, University of Pittsburgh

<sup>3</sup>Department of Psychiatry, Columbia University

<sup>4</sup>Department of Psychology, University of Oregon

{terrancel, pliang, morency}@cs.cmu.edu

October 17, 2020

## Abstract

Mental health conditions remain under-diagnosed even in countries with common access to advanced medical care. The ability to accurately and efficiently predict mood from easily collectible data has several important implications towards early detection and intervention of mental health disorders. One promising data source to help monitor human behavior is from daily smartphone usage. However, care must be taken to summarize behaviors without identifying the user through personal (e.g., personally identifiable information) or protected attributes (e.g., race, gender). In this paper, we study behavioral markers of daily mood using a recent dataset of mobile behaviors from high-risk adolescent populations. Using computational models, we find that multimodal modeling of both text and app usage features is highly predictive of daily mood over each modality alone. Furthermore, we evaluate approaches that reliably obfuscate user identity while remaining predictive of daily mood. By combining multimodal representations with privacy-preserving learning, we are able to push forward the performance-privacy frontier as compared to unimodal approaches.

## 1 Introduction

Mental illnesses can have a damaging permanent impact on communities, societies and economies all over the world. Individuals often do not realize they are at risk of mental disorders even when they have some mild symptoms. As a result, many are late in seeking professional help [30], particularly among adolescents where suicide is the second leading cause of death [5]. In addition to deaths, 16% of high school students report having seriously suicidal thoughts each year, and 8% of them make one or more suicide attempts [3]. Moreover, children with high irritability as well as depressive and anxious mood have a higher suicidal risk during adolescence compared with children with low symptoms [24]. As a step towards adaptive interventions of suicidal ideations, intensive monitoring of behaviors via adolescents' use of smartphones may shed new light on the early risk of suicidal thoughts and ideations [23]. While smartphones provide a valuable data source, one must take care to summarize behaviors without revealing user identities through personal (e.g., personally

---

\* first two authors contributed equally.

identifiable information) or protected attributes (e.g., race, gender) to potentially adversarial third-parties [26]. This form of anonymity is critical for implementing technologies in real-world scenarios.

Recent work in affective computing has begun to explore the potential in predicting mood and emotion from mobile data. In particular, they have found that typing patterns [2, 12, 15, 34], self-reporting apps [28], and wearable sensors [11, 13, 25] are particularly predictive for affect. In addition, multimodal modeling of multiple sensors (e.g. wearable sensors and smartphone apps) was shown to further improve performance [16, 29]. While current work primarily relies on self-report apps for *long-term* assessments of suicide risk [14], our work investigates fine-grained text typed and apps used throughout the entire day as a signal to detect *imminent* suicidal risk [9, 18], which is of critical clinical importance [14].

Recent studies have also shown that private traits are predictable from digital records of human behavior [17], which is dangerous especially when sensitive user data is concerned. As a result, in parallel to improving predictive performance, a recent focus has been on improving the privacy of real-world machine learning models through techniques such as differential privacy [6–8] and federated learning [22, 10, 20], especially for healthcare data (e.g. EHRs [33] and wearable devices [4]).

In this paper, as a step towards using *multimodal privacy-preserving* mood prediction as fine-grained signals to detect *imminent* suicidal risk, we analyze a recent dataset of mobile behaviors collected from adolescent populations at high suicidal risk. With consent from participating groups, the dataset collects fine-grained features spanning online communication, keystroke patterns, and application usage. Participants are administered clinical interviews probing daily mood scores. Using this unique dataset, we first find that machine learning algorithms on extracted word-level and app history features provide a promising learning signal in predicting daily self-reported mood scores. Furthermore, we find that joint multimodal modeling of both text and app usage features allows models to *contextualize* text typed in their corresponding apps which leads to better predictive performance over each modality alone. As a step towards privacy-preserving learning, we also evaluate approaches that obfuscate user identity while remaining predictive of daily mood. Finally, by combining multimodal contextualization with privacy-preserving learning, we are able to further push forward the performance-privacy frontier as compared to unimodal approaches.

## 2 Multimodal Mobile Dataset

Intensive monitoring of behaviors via adolescents’ frequent use of smartphones may shed new light on the early risk of suicidal thoughts and ideations [23]. Smartphones provide a valuable and natural data source with rich behavioral markers spanning online communication, keystroke patterns, and application usage. Learning these behavioral markers requires large datasets with diversity in participants, variety in features collected, and accuracy in annotations. As a step towards this goal, we have recently collected a dataset of mobile behaviors from high-risk adolescent populations with consent from participating groups.

We begin with a brief review of the data collection process. This data monitors adolescents spanning (a) recent suicide attempters (past 6 months) with current suicidal ideation, (b) suicide ideators with no past suicide attempts, and (c) psychiatric controls with no history of suicide ideation or attempts. Passive sensing data is collected from each participant’s smartphone across a duration of 6 months. Participants are administered clinical interviews probing for suicidal thoughts and behaviors (STBs) and self-report instruments regarding symptoms, and acute events (e.g., suicide attempts, psychiatric hospitalizations) are tracked weekly via a questionnaire. All users have given consent for their mobile device data to be collected and shared with us for research purposes.

**Labels:** Each day, users are asked to respond to the following question - “In general, how have you been feeling over the last day?” - with an integer score between 0 and 100, where 0 means very negative and 100 means very positive. To construct our prediction task, we discretized these scores into the following three bins - *negative* (0-33), *neutral* (34-66), *positive* (67-100), which follows a class distribution of 12.43%, 43.63%, and 43.94% respectively. For our 3-way classification task, participants with fewer than 50 daily self-reports were removed since such participants do not provide enough data to train an effective model. In total, our dataset consists of 1641 samples, consisting of data coming from 17 unique participants.

**Features:** For our study, we focused on keyboard data, which includes the time of data capture, the application used, and the text entered by the user. For each daily score response, we use keystroke information collected between 5AM on the previous day to 5AM on the current day. We then extracted the following features:

*Text:* After removing stop words, we collected the top 1000 words (out of approximately 3.2 million) used across all users in our dataset and created a *bag-of-words* feature that contains the daily number of occurrences of each word.

*App usage:* We count the number of keystrokes entered per each application, creating a *bag-of-apps* feature for each day. We discard applications that are used by less than 10% of the participants so that our features are generalizable to more than just a single user in the dataset, leaving 137 total apps (out of the original 640).

We observed that our models performed well when binarizing our feature vectors into boolean vectors, which signify whether a word or app was used on a given day (i.e. mapping values greater than 0 to 1). Our final feature vectors consist of a concatenation of a normalized and a binarized feature vector, resulting in 2000 and 274-dimensional vectors for Text and App features respectively.

### 3 Methods

In this paper we focus on studying approaches for learning privacy-preserving representations from mobile data for mood prediction. Our processed data comes in the form of triples  $\{(x_{t,i}, x_{a,i}, y_i)\}_{i=1}^n$  with  $x_t \in \mathbb{N}^{|V_t|}$  denoting the bag of text features with vocabulary size  $|V_t| = 2000$  and  $x_a \in \mathbb{N}^{|V_a|}$  denoting the bag of apps features with vocabulary size  $|V_a| = 274$ .  $y$  denotes the label which takes on 3 categories. In parallel, we also have data representing the corresponding (one-hot) user identity  $x_{id}$  which will be useful when learning privacy-preserving representations that do not encode information about  $x_{id}$  (and evaluating privacy performance).

**Unimodal approaches:** We begin with a description of 2 unimodal approaches we considered.

1. Support Vector Machines (SVMs) project training examples to a chosen kernel space and finds the optimal hyperplane that maximally separates each class of instances. We directly apply a SVM classifier on input data  $x_{uni} \in \{x_t, x_a\}$  and train with standard supervised learning to predict daily mood labels  $y$ .
2. Multilayer Perceptrons (MLPs) have seen widespread success in supervised prediction tasks due to their ability in modeling complex nonlinear relationships. Because of the small size of our dataset, we choose a simple multilayer perceptron with two hidden layers. Similarly, we apply an MLP classifier on input data  $x_{uni} \in \{x_t, x_a\}$  to predict daily mood labels  $y$ .

**Multimodal models:** We extend both SVM and MLP classifiers using early fusion [1] of text and app usage to model multimodal interactions. Specifically, we align the input through concatenating the bag-of-words

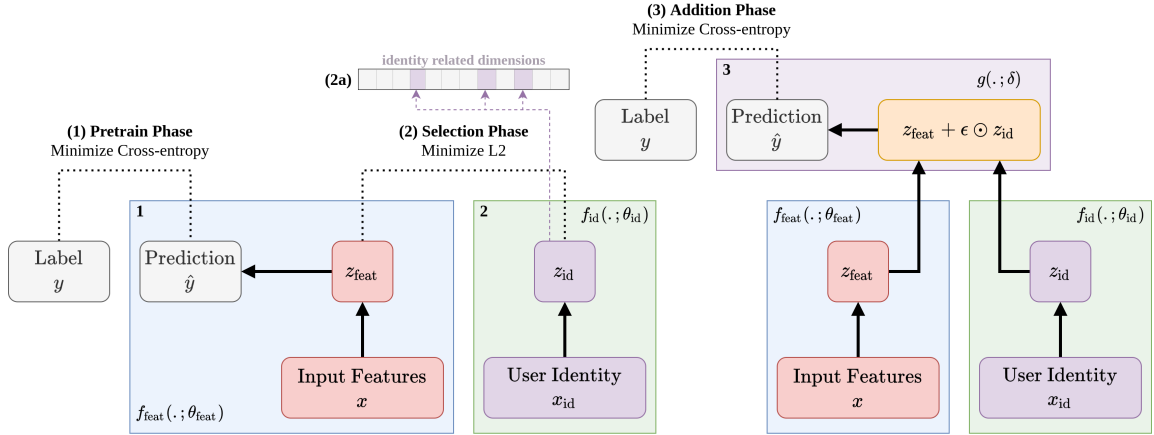


Figure 1: Diagram of the NI-MLP architecture. Boxes with numbers denote which parameters are being optimized in the corresponding step. For example in the addition phase (3), NI-MLP optimizes parameters  $\delta$  in  $g(\cdot; \delta)$ . 2a depicts  $z_{\text{id}}$ , which is a sparse vector of size  $\dim(z_{\text{feat}})$  whose nonzero values (colored purple) signify identity-confounding dimensions in  $z_{\text{feat}}$ .

and bag-of-apps features for each day before, resulting in an input vector  $x_{\text{multi}} = x_t \oplus x_a$ , before using a SVM/MLP classifier for prediction.

**A step toward preserving privacy:** While classifiers trained with supervised learning can learn useful representations for mood prediction, they carry the risk of *memorizing* the identity of the user along with their sensitive mobile usage and baseline mood scores, and possibly *revealing* these identities to adversarial third-parties. Therefore, we aim to achieve a balance between predictive performance of mood while protecting the privacy of personal identities.

We adapt the Selective-Additive Learning framework [31] with the goal of learning *disentangled* representations separated into *identity-dependent* and *identity-independent* factors using 2 phases:

*Selection phase:* We are first presented with a set of (multimodal) features that contain both identity-dependent and independent dimensions  $z_{\text{feat}} = f_{\text{feat}}(x; \theta_{\text{feat}}^*)$  by pretraining a MLP classifier using supervised learning.  $f_{\text{feat}}$  denotes the classifier with pretrained parameters  $\theta_{\text{feat}}^*$  and  $x$  denotes the input feature vector. Our goal is to now disentangle the identity-dependent and independent dimensions within  $z_{\text{feat}}$ , which can be done by training a separate classifier to predict  $z_{\text{feat}}$  as much as possible given only the user identity:

$$\theta_{\text{id}}^* = \arg \min_{\theta_{\text{id}}} (z_{\text{feat}} - f_{\text{id}}(x_{\text{id}}; \theta_{\text{id}}))^2 + \lambda \|\theta_{\text{feat}}\|_1, \quad (1)$$

where  $x_{\text{id}}$  denotes a one hot encoding of user identity,  $f_{\text{id}}$  denotes the identity encoder with parameters  $\theta_{\text{id}}$ , and  $\lambda$  denotes a hyperparameter that controls the weight of the L1 regularizer.  $f_{\text{id}}$  projects the user identity encodings to the feature space learned by  $f_{\text{feat}}$ . By minimizing the objective in equation 1 for each  $(x, x_{\text{id}})$  pair,  $f_{\text{id}}$  learns to encode user identity into a sparse vector  $z_{\text{id}} = f_{\text{id}}(x_{\text{id}}; \theta_{\text{id}}^*)$ , whose nonzero values represent identity-dependent dimensions in  $z_{\text{feat}}$ .

*Additive phase:* Given these 2 factors  $z_{\text{feat}}$  and  $z_{\text{id}}$  and to ensure that our prediction model does not capture identity-related information  $z_{\text{id}}$ , we add multiplicative Gaussian noise to the identity-related dimensions  $z_{\text{id}}$

while repeatedly optimizing for mood prediction with a final MLP classification layer  $g(z_{\text{feat}}, z_{\text{id}}; \delta)$ . This resulting model only retains identity-independent features for mood prediction:

$$\hat{y} = g(z_{\text{feat}} + \epsilon \odot z_{\text{id}}) \quad (2)$$

where  $\epsilon \sim \mathcal{N}(0, \sigma^2)$  is repeatedly sampled across batches and training epochs. We call this approach NOISY IDENTITY MLP, or NI-MLP for short.

**Controlling the tradeoff between performance and privacy:** We found that directly applying noise to the identity-dependent dimensions led to improved privacy at the expense of prediction performance. As a result, to control the tradeoff between performance and privacy, we vary the parameter  $\sigma$ , which controls amount of noise added to the identity-dependent dimensions across batches and training epochs.  $\sigma = 0$  recovers a standard MLP classifier with good performance but reveals user identities, while large  $\sigma$  effectively protects user identities but at the possible expense of mood prediction performance.

In practice, the optimal tradeoff between privacy and performance varies depending on the problem. For our purposes, we report in Table 1 and 2 results of models which achieve the highest performance-privacy ratio  $R$  on the validation set, where  $R = \frac{s_{\text{MLP}} - s_{\text{NI-MLP}}}{t_{\text{MLP}} - t_{\text{NI-MLP}}}$  is defined as the improvement in privacy per unit of performance lost. Here,  $s$  is defined as the accuracy in the user prediction task and  $t$  is defined as the F1 score on the mood prediction task.

## 4 Experiments

We perform experiments to test the utility of text and app features in predicting daily mood while keeping user identity private.

**Data splits:** Given that our data is longitudinal, we split our data into 10 partitions order chronologically by user. We do so in order to maintain independence between the train, validation, and test splits in the case where there is some form of time-level dependencies within our labels. For evaluating SVM and MLP, we run a nested k-fold cross validation (i.e. we perform 9-fold validation within 10-fold testing). For each test fold, we identify the optimal parameter set as the one that achieves highest mean validation score over the validation folds. We choose average F1 score as our validation metric because our classes are imbalanced. To evaluate NI-MLP, we use the best performing MLP model for each test fold as our base neural classifier.

**Results:** For our experiments we report the test accuracy and macro F1 score. Given the relatively low number of cross validation folds, use the Wilcoxon signed-rank test [32] at the 5% significance level for all significance testing. We make the following observations regarding our research questions.

**Observation 1: Text and app usage features are predictive of mood.** To evaluate how predictive bag-of-words and app usage features are, we run experiments using SVM, MLP, and NI-MLP. Because we have unbalanced classes, we choose a majority classifier (i.e. predict the most common class in the training set) for our baseline. Results are reported in Table 1. We observe that using our features, the models outperform the baseline with respect to accuracy and F1 score. Using the Wilcoxon test at the 5% significance level, we find that these improvements over the baseline in both F1 score and accuracy are statistically significant (p-value  $<< 0.05$ ).

**Observation 2: Fusing both text and app usage modalities improves performance.** In Table 1, we also compare the performance of our models on combined (Text + Apps) features to the performance on each individual feature set. We observe that for both metrics, combining the two features gives better performance over using either modality individually.

Table 1: Comparison of mood prediction performance across different modalities. Best results in **bold**. For both accuracy and F1 score, models trained on both text and apps features outperform models trained using individual modalities. \* and  $\diamond$  denote that the difference between multimodal and unimodal (Text and Apps respectively) is statistically significant ( $p$ -value  $\ll 0.05$ ).

Modalities	F1 score				Accuracy			
	BASELINE	SVM	MLP	NI-MLP	BASELINE	SVM	MLP	NI-MLP
Text + Apps	19.07	<b>62.08<math>^\diamond</math></b>	<b>58.38<math>^{*\diamond}</math></b>	<b>52.90</b>	40.18	<b>66.59<math>^\diamond</math></b>	<b>62.93<math>^{*\diamond}</math></b>	<b>56.76</b>
Text	19.07	61.15	56.27	52.63	40.18	65.83	60.61	56.08
Apps	19.07	58.65	52.29	51.32	40.18	62.65	55.26	55.68




Figure 2: Visualization of representations learned by (a) MLP and (b) NI-MLP, which have been reduced to two dimensions via t-SNE and colored by participant identity. Representations learned by NI-MLP are no longer separable by user which better preserve privacy.

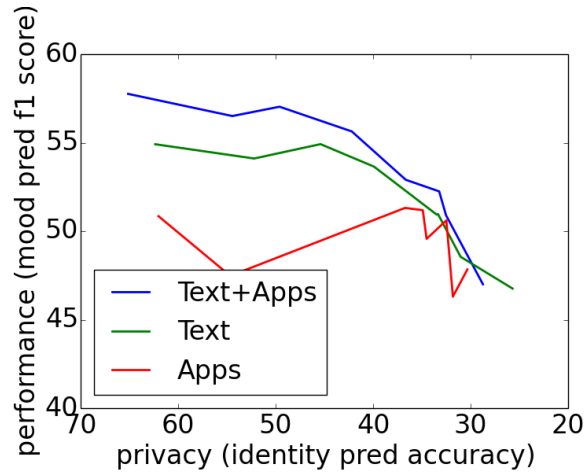
**Observation 3: Learning privacy-preserving features** While text and app history features are predictive of mood, we observe that such features are heavily correlated with user identity through both visual and empirical evaluations:

*Visual comparison.* We use t-SNE to reduce the learned features to 2 dimensions and after color-coding the points by participant identity, we see distinct clusters in Figure 2(a). After training NI-MLP, we again visualize the representations learned in Figure 2(b) and we find that they are no longer visually separable by user, indicating that NI-MLP indeed learns user-independent feature representations.

*Quantitative evaluation.* To empirically evaluate how well our models preserve privacy, we extracted each trained model’s penultimate layers and trained a logistic regression model to predict user identity from extracted feature representations. The more a model preserves privacy, the harder it should be to predict user identity. We observed that we can predict user identity based on the raw features with near perfect accuracy (95% using Text + Apps) which indicate that the labels are highly correlated with user identities. Next, we evaluated how each trained model’s feature representations encode information about user identities. While a logistic regression model can achieve up to 80% accuracy in predicting users from embeddings learned by the MLP, it achieves significantly lower accuracy when trained on NI-MLP embeddings - only 37%.

Table 2: **Left:** We evaluate privacy in predicting user identity from representations learned by MLP and NI-MLP (**lower** accuracy is better), and find that NI-MLP effectively obfuscates user identity. **Right:** Tradeoff between performance (mood prediction F1 score, **higher** is better) and privacy (identity prediction accuracy, **lower** is better) NI-MLP provides a tunable parameter  $\sigma$  which allows us to plot a range of (performance, privacy) points. Furthermore, using a multimodal model on text and app features pushes the Pareto front to obtain better performance and privacy at the same time.

	Text + Apps	Text	Apps
Raw features	95.70	92.65	91.82
MLP	79.04	76.41	85.94
NI-MLP	<b>36.65</b>	<b>38.38</b>	<b>36.72</b>



Note that we do not include results for SVM, which finds a linear separator in a specified kernel space rather than learning a representation for each sample. Moreover, explicitly projecting our features is computationally infeasible due to the dimensionality of our chosen kernel spaces.

*Understanding the tradeoff between performance and privacy.* NI-MLP provides a tunable parameter  $\sigma$  to control the privacy of user identities while remaining predictive of daily mood. In Table 2 (right) we plot the tradeoff between performance (mood prediction F1 score, higher is better) and privacy (identity prediction accuracy, lower is better). We find that using a multimodal model on both text and app features pushes the Pareto front to obtain better performance and privacy at the same time. Interestingly, we find that the tradeoff curve on a model trained only using app features does not exhibit a Pareto tradeoff curve as expected. We attribute this to randomness in predicting both mood and identities. Furthermore, Wang et al. [31] found that adding noise to identity dimensions can sometimes improve generalization by reducing reliance on identity-confounding dimensions, which could also explain occasional increased performance at larger  $\sigma$  values.

**Additional observations:** Since this is a new dataset, we explored several more machine learning methods throughout the research process. In this section we describe some of the approaches that yielded initial negative results despite them working for standard datasets in machine learning and natural language processing:

1. **Pretrained word and sentence representations:** We initially tried finetuning language model (BERT) and using word embedding (GLoVe) methods to extract text features, classical approaches which work extremely well for NLP problems. Surprisingly, we were not able to outperform classifiers based on bag-of-words features. We hypothesize that it is difficult to use these methods to capture such long sequences of data spread out over a single day. One avenue we wish to explore is to treat this feature extract as a ranking problem. This proposal is based on the assumption that in the course of a day, only a few key events/periods of time are predictive of mood, forcing us to design a method that identifies them.
2. **Keystroke features:** We also experimented with keystroke/typing pattern features by extracting several summary statistics as features (e.g. average number of words typed per minute), but these did not improve the model when combined with text and app features. Both the encoder for keystroke data as well as the better multimodal fusion with text and app features could help in this regard. For example,

fine-grained fusion that maps each keystroke to the application in which the text was entered in [21].

3. **User specific models:** We also explored the setting of training a separate model per user but we found that there was too little data per user to train a good model. As part of future work, we believe that if NI-MLP can learn a user independent neural classifier, this approach can be used pretrained representations for either finetuning or few shot learning on each specific user. Previous work in federated learning [20, 27] also offer ways of learning a user-specific model that leverages other users' data during training, which could help to alleviate the lack of data per user.
4. **User-independent data splits:** We have shown that text and app features are highly dependent on participant identities. Consequently, models trained on these features would perform poorly when evaluated on a user not found in the training set. We would like to evaluate if better learning of user independent features can improve generalization to new users (e.g. split the data such that the first 10 users are used for training, next 3 for validation, and final 4 for testing).

## 5 Conclusion

In this paper, we used multimodal machine learning to learn markers of daily mood from mobile data. Our method is also able to obfuscate user identity as a step towards more privacy-preserving learning, a direction crucial towards real-world learning from sensitive mobile data. Several important future directions involve: 1) investigating the relationship between user-dependent and independent features to characterize the tradeoffs between performance and privacy, 2) improving multimodal representations through more fine-grained temporal fusion [21], 3) incorporating more unique features provided in the dataset including keystrokes and geolocation data, and 4) federated training from decentralized data [22] to further improve privacy [10] and fairness [19] of sensitive mobile data.

## References

- [1] Tadas Baltrušaitis, Chaitanya Ahuja, and Louis-Philippe Morency. Multimodal machine learning: A survey and taxonomy. *IEEE transactions on pattern analysis and machine intelligence*, 41(2):423–443, 2018.
- [2] Bokai Cao, Lei Zheng, Chenwei Zhang, Philip S Yu, Andrea Piscitello, John Zulueta, Olu Ajilore, Kelly Ryan, and Alex D Leow. Deepmood: modeling mobile phone typing dynamics for mood detection. In *Proceedings of the 23rd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, pages 747–755, 2017.
- [3] CDC. *Suicide Facts at a Glance 2015*, 2015 (accessed September 6, 2020).
- [4] Yiqiang Chen, Xin Qin, Jindong Wang, Chaohui Yu, and Wen Gao. Fedhealth: A federated transfer learning framework for wearable healthcare. *IEEE Intelligent Systems*, 2020.
- [5] Sally C Curtin and Melanie P Heron. Death rates due to suicide and homicide among persons aged 10–24: United states, 2000–2017. 2019.
- [6] Fida Kamal Dankar and Khaled El Emam. The application of differential privacy to health data. In *Proceedings of the 2012 Joint EDBT/ICDT Workshops*, pages 158–166, 2012.
- [7] Fida Kamal Dankar and Khaled El Emam. Practicing differential privacy in health care: A review. *Trans. Data Priv.*, 6(1):35–67, 2013.



- [8] Fida Kamal Dankar, Khaled El Emam, Angelica Neisa, and Tyson Roffey. Estimating the re-identification risk of clinical data sets. *BMC medical informatics and decision making*, 12(1):66, 2012.
- [9] Joseph C Franklin, Jessica D Ribeiro, Kathryn R Fox, Kate H Bentley, Evan M Kleiman, Xieyining Huang, Katherine M Musacchio, Adam C Jaroszewski, Bernard P Chang, and Matthew K Nock. Risk factors for suicidal thoughts and behaviors: a meta-analysis of 50 years of research. *Psychological bulletin*, 143(2):187, 2017.
- [10] Robin C Geyer, Tassilo Klein, and Moin Nabi. Differentially private federated learning: A client level perspective. *arXiv preprint arXiv:1712.07557*, 2017.
- [11] Asma Ghandeharioun, Szymon Fedor, Lisa Sangermano, Dawn Ionescu, Jonathan Alpert, Chelsea Dale, David Sontag, and Rosalind Picard. Objective assessment of depressive symptoms with machine learning and wearable sensors data. In *2017 Seventh International Conference on Affective Computing and Intelligent Interaction (ACII)*, pages 325–332. IEEE, 2017.
- [12] Surjya Ghosh, Niloy Ganguly, Bivas Mitra, and Pradipta De. Evaluating effectiveness of smartphone typing as an indicator of user emotion. In *2017 Seventh International Conference on Affective Computing and Intelligent Interaction (ACII)*, pages 146–151. IEEE, 2017.
- [13] Surjya Ghosh, Niloy Ganguly, Bivas Mitra, and Pradipta De. Tapsense: Combining self-report patterns and typing characteristics for smartphone based emotion detection. In *Proceedings of the 19th International Conference on Human-Computer Interaction with Mobile Devices and Services*, pages 1–12, 2017.
- [14] Catherine R Glenn and Matthew K Nock. Improving the short-term prediction of suicidal behavior. *American journal of preventive medicine*, 47(3):S176–S180, 2014.
- [15] He Huang, Bokai Cao, S Yu Phillip, Chang-Dong Wang, and Alex D Leow. Dpmood: Exploiting local and periodic typing dynamics for personalized mood prediction. In *2018 IEEE International Conference on Data Mining (ICDM)*, pages 157–166. IEEE, 2018.
- [16] Natasha Jaques, Sara Taylor, Akane Sano, and Rosalind Picard. Multimodal autoencoder: A deep learning approach to filling in missing sensor data and enabling better mood prediction. In *2017 Seventh International Conference on Affective Computing and Intelligent Interaction (ACII)*, pages 202–208. IEEE, 2017.
- [17] Michal Kosinski, David Stillwell, and Thore Graepel. Private traits and attributes are predictable from digital records of human behavior. *Proceedings of the national academy of sciences*, 110(15):5802–5805, 2013.
- [18] Matthew Michael Large, Daniel Thomas Chung, Michael Davidson, Mark Weiser, and Christopher James Ryan. In-patient suicide: selection of people at risk, failure of protection and the possibility of causation. *BJPsych Open*, 3(3):102–105, 2017.
- [19] Paul Pu Liang, Irene Li, Emily Zheng, Yao Chong Lim, Ruslan Salakhutdinov, and Louis-Philippe Morency. Towards debiasing sentence representations. In *ACL*, 2020.
- [20] Paul Pu Liang, Terrance Liu, Liu Ziyin, Nicholas B. Allen, Randy P. Auerbach, David Brent, Ruslan Salakhutdinov, and Louis-Philippe Morency. Think locally, act globally: Federated learning with local and global representations. 2020.

- [21] Paul Pu Liang, Ziyin Liu, AmirAli Bagher Zadeh, and Louis-Philippe Morency. Multimodal language analysis with recurrent multistage fusion. In *EMNLP*, 2018.
- [22] H. Brendan McMahan, Eider Moore, Daniel Ramage, Seth Hampson, and Blaise Agüera y Arcas. Communication-efficient learning of deep networks from decentralized data. In *AISTATS*, 2016.
- [23] Inbal Nahum-Shani, Shawna N Smith, Bonnie J Spring, Linda M Collins, Katie Witkiewitz, Ambuj Tewari, and Susan A Murphy. Just-in-time adaptive interventions (jitais) in mobile health: key components and design principles for ongoing health behavior support. *Annals of Behavioral Medicine*, 52(6):446–462, 2018.
- [24] Massimiliano Orri, Cedric Galera, Gustavo Turecki, Alberto Forte, Johanne Renaud, Michel Boivin, Richard E Tremblay, Sylvana M Côté, and Marie-Claude Geoffroy. Association of childhood irritability and depressive/anxious mood profiles with adolescent suicidal ideation and attempts. *JAMA psychiatry*, 75(5):465–473, 2018.
- [25] Akane Sano, Sara Taylor, Andrew W McHill, Andrew JK Phillips, Laura K Barger, Elizabeth Klerman, and Rosalind Picard. Identifying objective physiological markers and modifiable behaviors for self-reported stress and mental health status using wearable sensors and mobile phones: observational study. *Journal of medical Internet research*, 20(6):e210, 2018.
- [26] Sagar Sharma, Keke Chen, and Amit Sheth. Toward practical privacy-preserving analytics for iot and cloud-based healthcare systems. *IEEE Internet Computing*, 22(2):42–51, 2018.
- [27] Virginia Smith, Chao-Kai Chiang, Maziar Sanjabi, and Ameet S Talwalkar. Federated multi-task learning. In *Advances in Neural Information Processing Systems*, pages 4424–4434, 2017.
- [28] Yoshihiko Suhara, Yinzhan Xu, and Alex’ Sandy’ Pentland. Deepmood: Forecasting depressed mood based on self-reported histories via recurrent neural networks. In *Proceedings of the 26th International Conference on World Wide Web*, pages 715–724, 2017.
- [29] Sara Ann Taylor, Natasha Jaques, Ehimwenma Nosakhare, Akane Sano, and Rosalind Picard. Personalized multitask learning for predicting tomorrow’s mood, stress, and health. *IEEE Transactions on Affective Computing*, 2017.
- [30] Graham Thornicroft, Nisha Mehta, Sarah Clement, Sara Evans-Lacko, Mary Doherty, Diana Rose, Mirja Koschorke, Rahul Shidhaye, Claire O’Reilly, and Claire Henderson. Evidence for effective interventions to reduce mental-health-related stigma and discrimination. *The Lancet*, 387(10023):1123–1132, 2016.
- [31] Haohan Wang, Aaksha Meghawati, Louis-Philippe Morency, and Eric P Xing. Select-additive learning: Improving generalization in multimodal sentiment analysis. In *2017 IEEE International Conference on Multimedia and Expo (ICME)*, pages 949–954. IEEE, 2017.
- [32] Frank Wilcoxon. Individual comparisons by ranking methods. In *Breakthroughs in statistics*, pages 196–202. Springer, 1992.
- [33] Jie Xu and Fei Wang. Federated learning for healthcare informatics. *arXiv preprint arXiv:1911.06270*, 2019.
- [34] John Zulueta, Andrea Piscitello, Mladen Rasic, Rebecca Easter, Pallavi Babu, Scott A Langenecker, Melvin McInnis, Olusola Ajilore, Peter C Nelson, Kelly Ryan, et al. Predicting mood disturbance severity with mobile phone keystroke metadata: a biaffect digital phenotyping study. *Journal of medical Internet research*, 20(7):e241, 2018.

Table 3: Model parameter configurations.  
 \*Integer kernel values denote the degree of a polynomial kernel.

Model	Parameter	Value
SVM	C	0.1, 0.5, 1, 2, 3, 5, 10
	Kernel*	RBF, 2, 3, 5, 10
MLP	hidden dim 1	1024, 512, 128
	hidden dim 2	128, 64
	dropout rate	0, 0.2, 0.5
	learning rate	0.001
	batch size	100
	epochs	200
NI-MLP	$\lambda$	0.1, 1, 2, 3, 5, 10
	$\sigma$	1, 5, 10, 25, 50, 100

## Appendix

### A Implementation Details

**Hyperparameters:** We show all hyperparameters used in Tables 3.