

476 A Motivation (Ext)

477 In this section, we delve deeper into the motivational aspects underscoring our research. Our study
478 primarily concerns tabular datasets where the total number of features is significantly outnumbered by
479 the training samples. We concentrate on applications necessitating user-provided personal information
480 for decision-making, such as lending, online insurance services, and health-care services. The
481 prevailing weaknesses of these applications, which our work attempts to address, are outlined as
482 follows:

- 483 1. **Privacy Concerns:** The primary issue stems from the need for users to disclose sensitive
484 information. For instance, in the context of online health-care services, patients are required
485 to share an array of sensitive health data—weight, height, smoking habits, etc.—via a
486 website or mobile application. This exposes users to potential privacy threats.
- 487 2. **User and Organizational Expenditure of Time and Effort:** Numerous applications, such
488 as lending, involve the time-consuming and effort-intensive process of gathering sensitive
489 information and its supporting evidence. Users, for example, are required to validate
490 their income through paystips or employment contracts in lending applications. Similarly,
491 the organization must invest time and resources to verify the authenticity of submitted
492 documents.
- 493 3. **Legal Constraints:** According to the EU General Data Protection Regulation’s principle
494 of data minimization, the collection of excessive personal information by companies and
495 organizations is restricted. Our primary text illustrates that it is not imperative to report all
496 features to preserve the model’s accuracy.

497 **The Imperative for Minimal Inference Time** We also emphasize the importance of low inference
498 time from both user and business perspectives. From the user’s viewpoint, applications such as online
499 car insurance require answering a series of questions to determine the insurance plan. Naturally, users
500 prefer answering fewer questions in the least amount of time. From a business standpoint, prolonged
501 inference time may lead to customer dissatisfaction, potentially resulting in contract termination.
502 This serves as the basis for our algorithmic choices, in lieu of more complex conditional distribution
503 modeling methods, which can significantly increase inference time.

504 B Related work

505 While we are not aware of studies on data minimization for inference problems, we draw connections
506 with differential privacy, feature selection, and active learning.

507 **Differential Privacy.** Differential Privacy (DP) [7] is a strong privacy notion which determines and
508 bounds the risk of disclosing sensitive information of individuals participating into a computation.
509 In the context of machine learning, DP ensures that algorithms can learn the relations between data
510 and predictions while preventing them from memorizing sensitive information about any specific
511 individual in the training data. In such a context, DP is primarily adopted to protect training data
512 [1, 6, 24] and thus the setting contrasts with that studied in this work, which focuses on identifying the
513 superfluous features revealed by users at *test time* to attain high accuracy. Furthermore, achieving tight
514 constraints in differential privacy often comes at the cost of sacrificing accuracy, while the proposed
515 privacy framework can reduce privacy loss without sacrificing accuracy under the assumption of
516 linear classifiers.

517 **Feature selection.** Feature selection [5] is the process of identifying and selecting a relevant subset
518 of features from a larger set for use in model construction, with the goal of improving performance
519 by reducing complexity and dimensionality of the data. The problem studied in this work can be
520 considered as a specialized form of feature selection with the added consideration of personalized
521 levels, where each individual may use a different subset of features. This contrasts standard feature
522 selection [13], which select the same subset of features for each data sample. Additionally, and unlike
523 traditional feature selection, which is performed during training and independent of the deployed
524 classifier [5], the proposed framework performs feature selection at deployment time and is inherently
525 dependent on the deployed classifier.

526 **Active learning.** Finally, the proposed framework shares similarities with active learning [8, 20],
527 whose goal is to iteratively select samples for experts to label in order to construct an accurate

528 classifier with the least number of labeled samples. Similarly, the proposed framework iteratively
 529 asks individuals to reveal one attribute given their released features so far, with the goal of minimizing
 530 the uncertainty in model predictions.

531 Despite these similarities, the proposed data minimization for inference concept is motivated by a
 532 privacy need and pertains to the analysis of features to release to induce the same level of accuracy as
 533 if all features were released.

534 C Missing proofs

535 **Proposition 1.** *Given a core feature set $R \subseteq S$ with failure probability $\delta < 0.5$, then there exists a*
 536 *function $\epsilon : \mathbb{R} \rightarrow \mathbb{R}$ that is monotonic decreasing function with $\epsilon(1) = 0$ such that:*

$$H[f_\theta(X_U, X_R = x_R)] \leq \epsilon(1 - \delta),$$

537 where $H[Z] = -\sum_{z \in [L]} \Pr(Z = z) \log \Pr(Z = z)$ is the entropy of the random variable Z .

538 *Proof.* In this proof, we demonstrate the binary classification case. The extension to a multi-class
 539 scenario can be achieved through a similar process.

540 By the definition of the core feature set, there exists a representative label, denoted as $\tilde{y} \in \{0, 1\}$
 541 such that the probability of $P(f_\theta(X_U, X_R = x_R) = \tilde{y})$ is greater than or equal to $1 - \delta$. Without
 542 loss of generality, we assume that the representative label is $\tilde{y} = 1$. Therefore, if we denote Z as
 543 the probability of $Pr(f_\theta(X_U, X_R = x_R) = 1)$, then the probability of $Pr(f_\theta(X_U, X_R = x_R) =$
 544 $0) = 1 - Z$. Additionally, we have $Z \geq 1 - \delta > 0.5$ due to the definition of core feature set
 545 and by the assumption that $\delta < 0.5$. The entropy of the model's prediction can be represented as:
 546 $H[f_\theta(X_U, X_R = x_R)] = -Z \log Z - (1 - Z) \log(1 - Z)$.

547 Choose $\epsilon(Z) = -Z \log Z - (1 - Z) \log(1 - Z)$. The derivative of $\epsilon(Z)$ is given by $\frac{d\epsilon(Z)}{dZ} =$
 548 $\log \frac{1-Z}{Z} < 0$, as $Z > 0.5$. As a result, $\epsilon(Z)$ is a monotonically decreasing function, so $\epsilon(Z) \leq$
 549 $\epsilon(1 - \delta)$

550 When $\delta = 0$, we have $Z = 1$, and by the property of the entropy $H[f_\theta(X_U, X_R = x_R)] = 0$. \square

551 **Proposition 2.** *Given two subsets R and R' of sensitive features S , with $R \subseteq R'$,*

$$H(f_\theta(X_U, X_R = x_R)) \geq H(f_\theta(X_{U'}, X_{R'} = x_{R'})),$$

552 where $U = S \setminus R$ and $U' = S \setminus R'$.

553 *Proof.* This is due to the property that conditioning reduces the uncertainty, or the well-known
 554 *information never hurts* theorem in information theory [9]. \square

555 **Proposition 3.** *The conditional distribution of any subset of unrevealed features $U' \in U$, given the*
 556 *the values of released features $X_R = x_R$ is given by:*

$$\Pr(X_{U'} | X_R = x_R) = \mathcal{N}\left(\mu_{U'} + \Sigma_{U',R} \Sigma_{RR}^{-1} (x_R - \mu_R), \Sigma_{U'U'} - \Sigma_{U'R} \Sigma_{RR}^{-1} \Sigma_{R,U'}\right),$$

557 where Σ is the covariance matrix

558 *Proof.* This is a well-known property of the Gaussian distribution and we refer the reader to Chapter
 559 2.3.2 of the textbook [3] for further details. \square

560 **Proposition 4.** *The model predictions before thresholding, $\tilde{f}_\theta(X_U, X_R = x_R) = \theta_U X_U + \theta_R x_R$ is*
 561 *a random variable with a Gaussian distribution $\mathcal{N}(m_f, \sigma_f)$, where*

$$m_f = \theta_R x_R + \theta_U^\top (\mu_U + \Sigma_{UR} \Sigma_{RR}^{-1} (x_R - \mu_R)) \quad (8)$$

$$\sigma_f^2 = \theta_U^\top (\Sigma_{UU} - \Sigma_{UR} \Sigma_{RR}^{-1} \Sigma_{RU}) \theta_U, \quad (9)$$

562 where θ_U is the sub-vector of parameters θ corresponding to the unrevealed features U .

563 *Proof.* The proof of this statement is straightforward due to the property that a linear combination of
564 Gaussian variables X_U is also Gaussian. Additionally, the posterior distribution of X_U is already
565 provided in Proposition 3. \square

566 **Proposition 5.** *Let the model predictions prior thresholding $\tilde{f}_\theta(X_U, X_R = x_R)$, be a random*
567 *variable following a Gaussian distribution $\mathcal{N}(m_f, \sigma_f^2)$. Then, the model prediction following thresh-*
568 *olding $f_\theta(X_U, X_R = x_R)$ is a random variable following a Bernoulli distribution $Bern(p)$ with*
569 *$p = \Phi(\frac{m_f}{\sigma_f})$, where $\Phi(\cdot)$ refers to the CDF of the standard normal distribution, and m_f and σ_f , are*
570 *given in Equations (5) and (6), respectively.*

571 *Proof.* In the case of a binary classifier, we have $f_\theta(x) = \mathbf{1}\{\tilde{f}_\theta(x) \geq 0\}$. If \tilde{f} follows a normal
572 distribution, denoted as $\tilde{f} \sim \mathcal{N}(m_f, \sigma_f^2)$, then by the properties of the normal distribution, f_θ follows
573 a Bernoulli distribution, denoted as $f_\theta \sim Bern(p)$, with parameter $p = \Phi(\frac{m_f}{\sigma_f})$, where $\Phi(\cdot)$ is the
574 cumulative density function of the standard normal distribution. \square

575 **Proposition 6.** *Assume f_θ is a linear classifier. Then, determining if a subset U of sensitive features*
576 *S is a pure core feature set can be performed in $O(|P| + |S|)$ time.*

577 *Proof.* As discussed in the main text, to test if a subset U is a core feature set or not, we need to
578 check if the following two terms have the same sign (either negative or non-negative):

$$\begin{aligned} \max_{X_U} \theta_U^\top X_U + \theta_R^\top x_R &= \|\theta_U\|_1 + \theta_R^\top x_R \\ \min_{X_U} \theta_U^\top X_U + \theta_R^\top x_R &= -\|\theta_U\|_1 + \theta_R^\top x_R. \end{aligned} \quad (10)$$

579 These can be solved in time $O(|P| + |S|)$ due to the property of the linear equality above. \square

580 **Theorem 1.** *The distribution of the random variable $\tilde{f}_\theta = \tilde{f}_\theta(X_U, X_R = x_R)$ where $X_U \sim$*
581 *$\mathcal{N}(\mu_U^{pos}, \Sigma_U^{pos})$ can be approximated by a Normal distribution as*

$$\tilde{f}_\theta \sim \mathcal{N}(\tilde{f}_\theta(X_U = \mu_U^{pos}, X_R = x_R), g_U^\top \Sigma_U^{pos} g_U) \quad (11)$$

582 where $g_U = \nabla_{X_U} \tilde{f}_\theta(X_U = \mu_U^{pos}, X_R = x_R)$ is the gradient of model prediction at $X_U = \mu_U^{pos}$.

583 *Proof.* The proof relies on the first Taylor approximation of classifier \tilde{f} around its mean:

$$\tilde{f}_\theta(X_U, X_R = x_R) \approx \tilde{f}_\theta(X_U = \mu_U^{pos}, X_R = x_R) + (X_U - \mu_U^{pos})^T \nabla_{X_U} \tilde{f}_\theta(X_U = \mu_U^{pos}, X_R = x_R) \quad (12)$$

584 Since $X_U \sim \mathcal{N}(\mu_U^{pos}, \Sigma_U^{pos})$ hence $X_U - \mu_U^{pos} \sim \mathcal{N}(\mathbf{0}, \Sigma_U^{pos})$. By the properties of normal
585 distribution, the right-hand side of Equation (12) is a linear combination of Gaussian variables, and it
586 is also Gaussian. \square

587 D Algorithms Pseudocode

588 The pseudocode for MinDRel for non-linear classifiers is presented in Algorithm 2. There are two
589 main differences between this algorithm and the case of linear classifiers. Firstly, unlike linear
590 classifiers, the procedure of pure core feature testing on line 5 does not require the guarantee (see
591 again Section 6.2). The accuracy of the testing procedures depends on the number of random samples
592 that we evaluate. The greater the number of drawn samples, the more likely the testing procedure
593 is to be accurate. During experiments, we draw 10^5 samples to perform the testing. Additionally,
594 we use Theorem 1 to estimate the distribution of the soft prediction as seen on line 11, as the exact
595 distribution cannot be computed analytically as in the case of linear classifiers.

596 E Extension from binary to multiclass classification

597 In the main text, we provide the implementation of MinDRel for binary classification problem. In
598 this section, we extend the method to the multiclass classification problem.

Algorithm 2: MinDRel for non-linear classifiers

input : A test sample x ; training data D
output : A core feature set R and its representative label \tilde{y}

```
1  $\mu \leftarrow \frac{1}{|D|} \sum_{(x,y) \in D} x$ 
2  $\Sigma \leftarrow \frac{1}{|D|} \sum_{(x,y) \in D} (x - \mu)(x - \mu)^\top$ 
3 Initialize  $R = \emptyset$ 
4 while True do
5   if  $R$  is a core feature set with repr. label  $\tilde{y}$  then
6     return  $(R, \tilde{y})$ 
7   else
8     foreach  $j \in U$  do
9       Compute  $\Pr(X_j | X_R = x_R)$  (using Prop. 3)
10       $Z \leftarrow \text{sample}(\Pr(X_j | X_R = x_R))$   $T$  times
11      Compute  $\Pr(f_\theta(X_j = z, X_{U \setminus \{j\}} | X_R = x_R))$  using Theorem 1)
12      Compute  $F(X_j)$  (using Eq. (4))
13     $j^* \leftarrow \text{argmax}_j F(X_j)$ 
14     $R \leftarrow R \cup \{j^*\}$ 
15     $U \leftarrow U \setminus \{j^*\}$ 
```

599 **E.1 Estimating $P(f_\theta(X_U, X_R = x_R))$**

600 In order to achieve our goals of determining if a subset is a core feature set for a given $\delta > 0$, and
601 computing the entropy in the scoring function, we need to estimate the distribution of $f_\theta(X_U, X_R =$
602 $x_R)$. In this section, we first discuss the method of computing the distribution of $\tilde{f}_\theta(X_U, X_R = x_R)$
603 for both linear and non-linear models. Once this is done, we then address the challenge of estimating
604 the hard label distribution $P(f_\theta(X_U, X_R = x_R))$.

605 It is important to note that, under the assumption that the input features X are normally distributed
606 with mean μ and covariance matrix Σ , the linear classifier $\tilde{f}_\theta = \theta^\top x$ will also have a multivariate
607 normal distribution. Specifically, if $X_U \sim \mathcal{N}(\mu_U^{pos}, \Sigma_U^{pos})$, then $\tilde{f}_\theta(X_U, X_R = x_R) \sim \mathcal{N}(\theta_R^\top x_R +$
608 $\theta_U^\top \mu_U^{pos}, \theta_U^\top \Sigma_U \theta_U)$.

609 For non-linear classifiers, the output $f_\theta(X_U, X_R = x_R)$ is not a Gaussian distribution due to the
610 non-linear transformation. To approximate it, we use Theorem 1 which states that the non-linear
611 function $\tilde{f}_\theta(X_U, X_R = x_R)$ can be approximated as a multivariate Gaussian distribution.

612 **Challenges when estimating $P(f_\theta(X_U, X_R = x_R))$** For multi-class classification problems, the
613 hard label $f_\theta(X_U, X_R = x_R)$ is obtained by selecting the class with the highest score, which
614 is given by $\text{argmax}_{i \in [L]} \tilde{f}_\theta^i(X_U, X_R = x_R)$. However, due to the non-analytical nature of the
615 argmax function, even when $\tilde{f}_\theta(X_U, X_R = x_R)$ follows a Gaussian distribution, the distribution
616 of $f_\theta(X_U, X_R = x_R)$ cannot be computed analytically. To estimate this distribution, we resort to
617 Monte Carlo sampling. Specifically, we draw a number of samples from $P(\tilde{f}_\theta(X_U, X_R = x_R))$, and
618 for each class $y \in \mathcal{Y}$ we approximate the probability $P(f_\theta(X_U, X_R = x_R) = y)$ as the proportion
619 of samples that fall in that class y .

620 We provide experiments of MinDRel for multi-class classification cases in Section F.5.

621 **F Experiments details**

622 **Datasets information** To show the advantages of the suggested MinDRel technique for safeguard-
623 ing feature-level privacy, we employ benchmark datasets in our experiments. These datasets include
624 both binary and multi-class classification datasets. The following are examples of binary datasets that
625 we use to evaluate the method:

- 626 1. Bank dataset [4]. The objective of this task is to predict whether a customer will subscribe
627 to a term deposit using data from various features, including but not limited to call duration
628 and age. There are a total of 16 features available for this analysis.
- 629 2. Adult income dataset [4]. The goal of this task is to predict whether an individual earns
630 more than \$50,000 annually. After preprocessing the data, there are a total of 40 features
631 available for analysis, including but not limited to occupation, gender, race, and age.
- 632 3. Credit card default dataset [4]. The objective of this task is to predict whether a customer
633 will default on a loan. The data used for this analysis includes 22 different features, such as
634 the customer’s age, marital status, and payment history.
- 635 4. Car insurance dataset [19]. The task at hand is to predict whether a customer has filed a
636 claim with their car insurance company. The dataset for this analysis is provided by the
637 insurance company and includes 16 features related to the customer, such as their gender,
638 driving experience, age, and credit score.

639 Furthermore, we also evaluate our method on two additional multi-class classification datasets:

- 640 1. Customer segmentation dataset [22]. The task at hand is to classify a customer into one of
641 four distinct categories: A, B, C, and D. The dataset used for this task contains 9 different
642 features, including profession, gender, and working experience, among others.
- 643 2. Children fetal health dataset [12]. The task at hand is to classify the health of a fetus into
644 one of three categories: normal, suspect, or pathological, using data from CTG (cardiotocog-
645 raphy) recordings. The data includes approximately 21 different features, such as heart rate
646 and the number of uterine contractions.

647 **Settings:** For each dataset, 70% of the data will be used for training the classifiers, while the
648 remaining 30% will be used for testing. The number of sensitive features, denoted as $|S|$, will be
649 chosen randomly from the set of all features. The remaining features will be considered as public.
650 100 repetition experiments will be performed for each choice of $|S|$, under different random seeds,
651 and the results will be averaged. All methods that require Monte Carlo sampling will use 100 random
652 samples. The performance of different methods will be evaluated based on accuracy and data leakage.
653 Two different classifiers will be considered.

- 654 1. Linear classifiers: We use Logistic Regression as the base classifier.
- 655 2. Nonlinear classifiers: The nonlinear classifiers used in this study consist of a neural network
656 with two hidden layers, using the ReLU activation function. The number of nodes in each
657 hidden layer is set to 10. The network is trained using stochastic gradient descent (SGD)
658 with a batch size of 32 and a learning rate of 0.001 for 300 epochs.

659 For Bayesian NN, we employ the package *bayesian-torch* [10] with the default settings. The base
660 regressor is a neural network with one hidden layer that has 10 hidden nodes and a ReLU activation
661 function. We train the network in 300 epochs with learning rate of 0.001.

662 **Baseline models.** We compare our proposed algorithms with the following baseline models:

- 663 1. **All features:** This refers to the usage of original classifier which asks users to reveal **all**
664 sensitive features.
- 665 2. **Optimal:** This method involves evaluating all possible subsets of sensitive features ($2^{|S|}$ in
666 total) in order to identify the minimum *pure* core feature set. For each subset, the verification
667 algorithm is used to determine whether it is a pure core feature set. The minimum pure
668 core feature set that is found is then selected. It should be noted that as all possible subsets
669 are evaluated, all sensitive feature values must be revealed. Therefore, this approach is not
670 practical in real-world scenarios. However, it does provide a lower bound on data leakage
671 for MinDRel (when $\delta = 0$).

672 **MinDRel models** In MinDRel there are two important steps: (1) core feature set verification and
673 (2) selection next feature to reveal. As additional baselines, we keep the core feature set verification
674 and vary the selection process. We consider the following three feature selection methods:

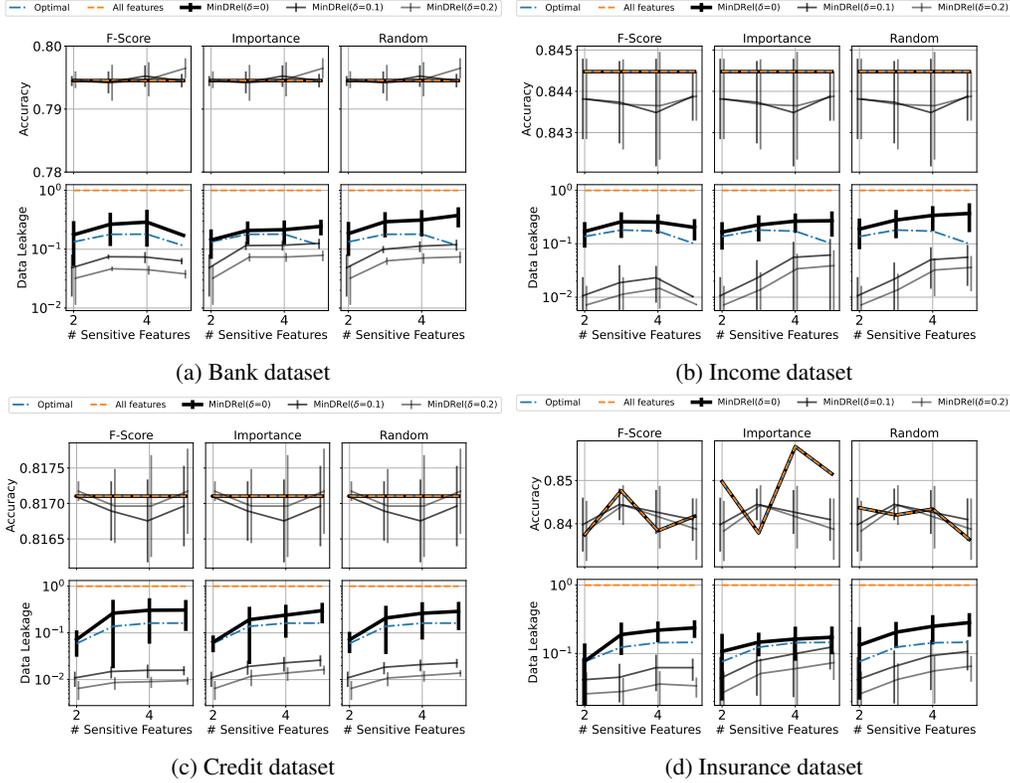


Figure 7: Comparison between using (left) our proposed F-Score (left) with Importance (Middle) and Random (Right) for different choices of number of sensitive features $|S|$. The baseline classifier is Logistic Regression

- 675 1. **F-Score:** We choose the feature based on amount of information on model prediction we
676 gain after revealing one feature as provided in Equation 3.
- 677 2. **Importance:** We reveal the unknown sensitive features based on the descending order of
678 feature importance until we find a core feature set. The feature importance is determined as
679 follows. We firstly fit a Logistic Regression $f_{\theta}(x) = 1\{\theta^T x \geq 0\}$ on the training dataset
680 D using all features (public included). The importance of one sensitive feature $i \in S$
681 is determined by $\|\theta_i\|_2$.
- 682 3. **Random:** We reveal the unrevealed sensitive feature in a random order until the revealed set
683 is a core feature set.

684 **Metrics.** We compare all different algorithms in terms of accuracy and data leakage:

- 685 1. Accuracy. For algorithms that are based on the core feature set, such as our MinDRel and
686 Optimal, the representative label is used as the model's prediction. Again, the representative
687 label for $\delta = 0$ can be identified by using testing pure core feature set procedures. For $\delta > 0$,
688 the representative label is given by $\tilde{y} = \operatorname{argmax}_{y \in \mathcal{Y}} \int P(f_{\theta}(X_U = x_U, X_R = x_R) =$
689 $y)P(X_U|X_R = x_R)dx_u$. The accuracy is then determined by comparing this representative
690 label to the ground truth.
- 691 2. Data leakage. We compute the percentage of the number sensitive features that users need
692 to provide on the test set. A small data leakage is considered better.

693 F.1 Additional comparison between using Gaussian assumption and Bayesian NN

694 We first show empirically the benefits of our proposed Gaussian assumption compared to using
695 Bayesian NN which allows more flexibility in modeling the conditional distribution $P(X_U|X_R = x_R)$.

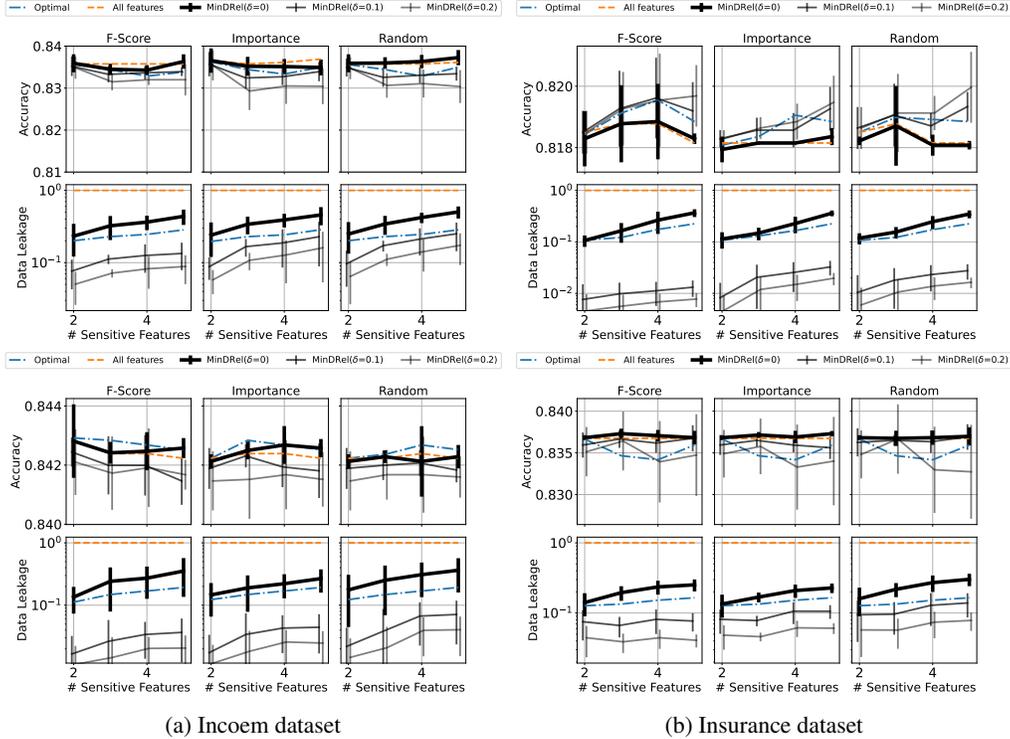


Figure 8: Comparison between using (left) our proposed F-Score (left) with Importance (Middle) and Random (Right) for different choices of number of sensitive features $|S|$. The baseline classifier is a neural network classifier.

Table 1: Comparison between using Bayesian neural network and our Gaussian assumption in term of training time (minutes) when $|S| = 5$ for various datasets.

Method	Bank	Income	Credit	Insurance
Bayesian NN	204	375	125	90
Gaussian assumption	0.01	0.02	0.02	0.01

696 We report both training and inference time between Bayesian NN and our Gaussian assumption on
697 various datasets when the number of sensitive features $|S| = 5$ in Table 1 and Table 2. When $|S| = 5$
698 the number of possible subsets $U \in S$ is $2^5 = 32$ which requires training 32 Bayesian NN models.
699 This will be especially slow for datasets with large number of training samples (e.g., Income with 50K
700 samples). In contrast, using Gaussian assumption we just need to precompute 32 inverse matrices
701 $\Sigma_{R,R}^{-1}$ which is pretty fast for data that have small number features (less than 50 in our experiments).
702 It is noted again that in this paper we focus on the case when the number of training samples is much
703 more than number of features. Likewise, during inference time, with Gaussian assumption we can
704 compute the distribution of model prediction in a closed form by simple matrix multiplication which
705 takes $O(d^2)$. Instead, using Bayesian NN, it requires expensive Monte Carlo sampling, especially
706 when $|U|$ is large to obtain an accurate estimation of $P(X_U|X_R = x_R)$.

707 We also report the performance in term of accuracy and data leakage between using Gaussian
708 assumption and Bayesian NN in Figure 9. We see no much significant difference in term of accuracy
709 and data leakage between two choices of modeling $P(X_U|X_R = x_R)$. In addition, as indicated above
710 using Gaussian assumption reduces significantly the training and inference time, in the subsequent
711 experiments we will use the Gaussian assumption in MinRDel with F-Score selection.

Table 2: Comparison between using Bayesian neural network and our Gaussian assumption in term of inference time (minutes) on test set when $|S| = 5$, $\delta = 0$ for various datasets.

Method	Bank	Income	Credit	Insurance
Bayesian NN	40	254	220	34
Gaussian assumption	15	78	66	9

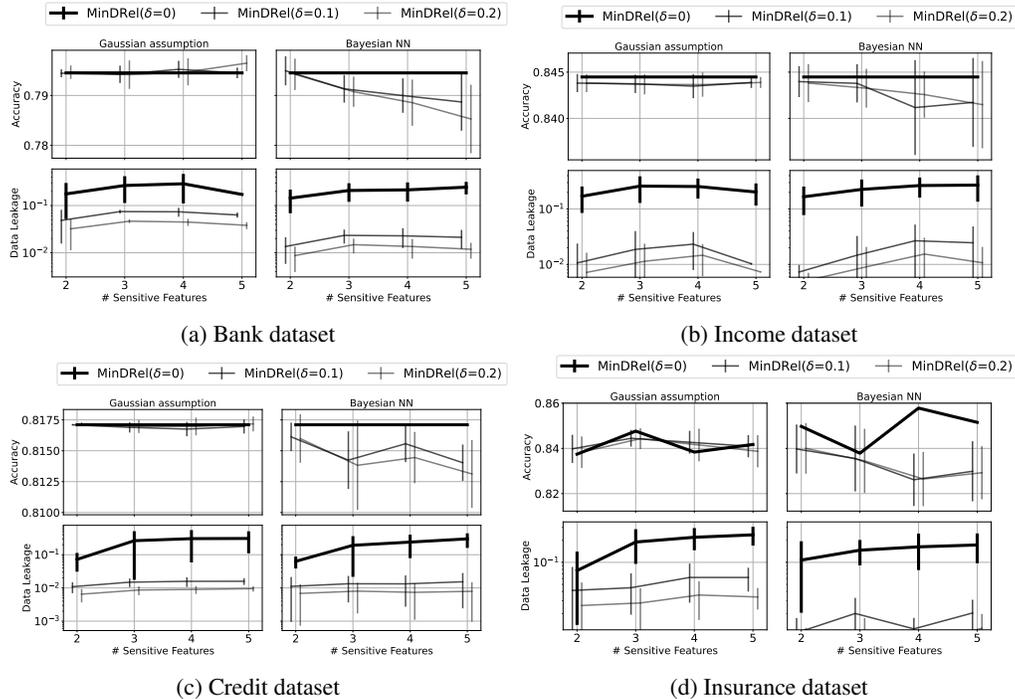


Figure 9: Comparison between using Bayesian NN with our Gaussian assumption in term of (1): accuracy and (2) data leakage for different choices of number of sensitive features $|S|$ on different datasets using a Logistic Regression classifier.

712 F.2 Additional experiments on linear binary classifiers

713 Additional experiments were conducted to compare the performance of MinDRel to that of the
 714 baseline methods using linear classifiers on the Bank, Adult income, Credit and Insurance datasets, as
 715 shown in Figure 7. As in the main text, a consistent trend in terms of performance is observed. As the
 716 number of sensitive attributes, $|S|$, increases, the data leakage introduced by MinDRel with various
 717 values of δ increases at a slower rate. With different choices of $|S|$, MinDRel (with $\delta = 0$) requires
 718 the revelation of at most 50% of sensitive information. To significantly reduce the data leakage of
 719 MinDRel, the value of δ can be relaxed. As mentioned in the main text, δ controls the trade-off
 720 between accuracy and data leakage here. The larger δ is, the greater uncertainty the model prediction
 721 has, which implies the fewer number of sensitive features users need to reveal and the lower accuracy
 722 on model prediction. By choosing an appropriate value for the failure probability, such as $\delta = 0.1$,
 723 only minimal accuracy is sacrificed (at most 0.002%), while the data leakage can be reduced to as
 724 low as 5% of the total number of sensitive attributes.

725 F.3 Additional experiments on non-linear binary classifiers

726 Additional experiments were conducted to compare the performance of MinDRel to that of the
 727 baseline methods using non-linear classifiers on the Bank, Adult income, Credit and Insurance
 728 datasets, as shown in Figure 8. As seen, while the baseline **All features** method requires the
 729 revelation of all sensitive attributes, MinDRel with different values of δ only requires the revelation of
 730 a much smaller number of sensitive attributes. The accuracy difference between the Baseline method

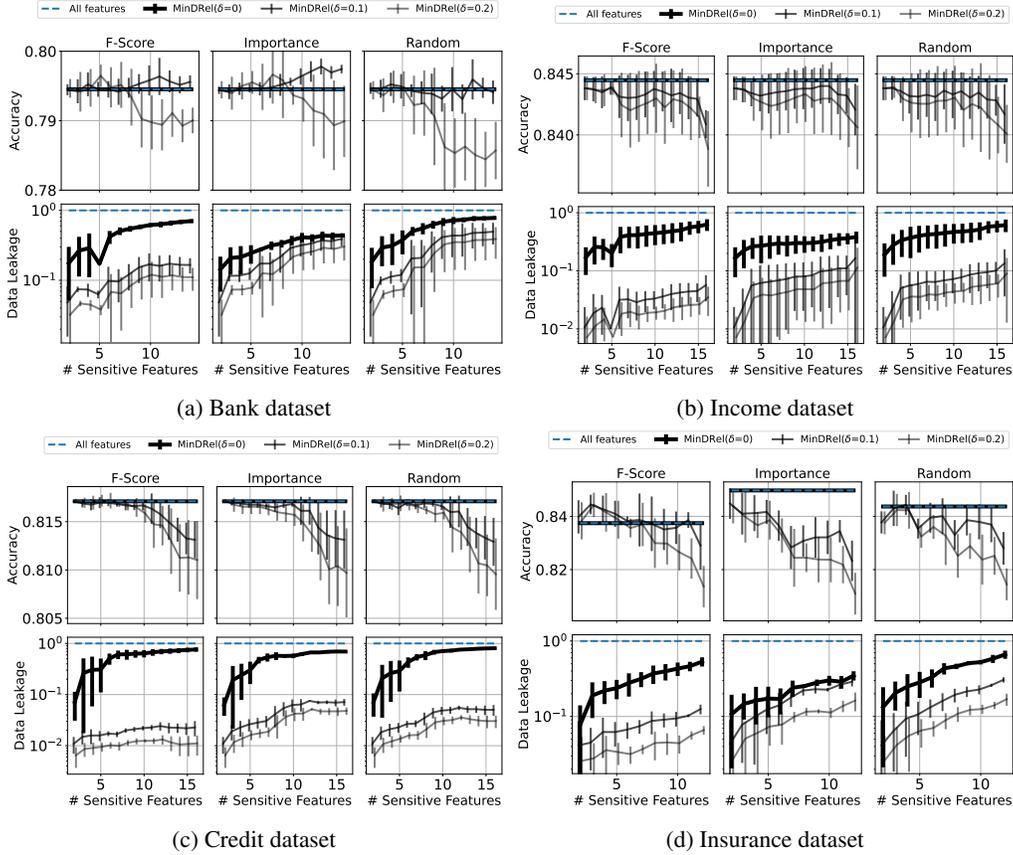


Figure 10: Comparison between using (left) our proposed F-Score (left) with Importance (Middle) and Random (Right) for different choices of number of sensitive features $|S|$. The baseline classifier is a logistic regression classifier.

731 and MinDRel is also minimal (at most 2%). These results demonstrate the effectiveness of MinDRel
 732 in protecting privacy while maintaining a good prediction performance for test data.

733 F.4 Scalability of MinDRel for large $|S|$

734 We demonstrate the performance of MinDRel when we have a large number of sensitive features $|S|$.
 735 Note that to reduce the runtime we did not run *Optimal* method which performs an exponential search
 736 over all possible choices of subset of S .

737 We first report the accuracy and data leakage of MinDRel when using F-Score or using either two
 738 heuristic rules Importance and Random in case of logistic regression classifiers in Figure 10.

739 Finally, we report the average testing time (in seconds) to get the model prediction per user of
 740 MinDRel in Figure 11. It is noted that in this case, we assume the time taken by users to release
 741 sensitive features is negligible. It is evident that when when $|S| > 15$, our proposed MinDRel with
 742 F-Score can take slightly more than 1 second to get the model prediction per user. This demonstrates
 743 the applicability of the models in practices.

744 F.5 Evaluation of MinDRel on multi-class classifiers

745 **Linear classifiers** We also provide a comparison of accuracy and data leakage between our proposed
 746 MinDRel and the baseline models for linear classifiers. These metrics are reported for the Customer
 747 and Children Fetal Health datasets in Figures 12a and 12b, respectively. The figures clearly shows
 748 the benefits of MinDRel in reducing data leakage while maintaining a comparable accuracy to the
 749 baseline models.

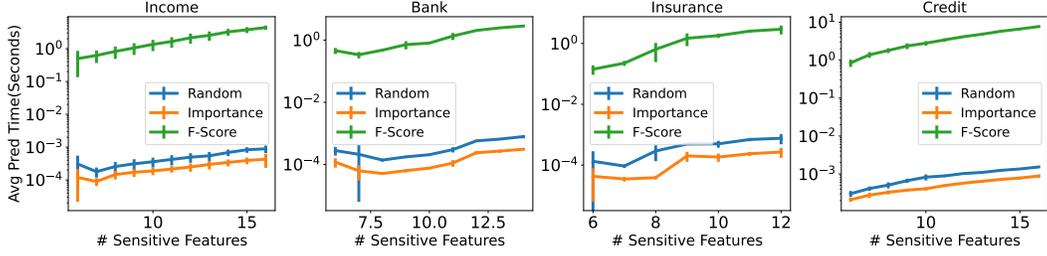


Figure 11: Comparison in term of average prediction time (seconds) among F-Score, Importance and Random method in MinDRel ($\delta = 0$) for different $|S|$.

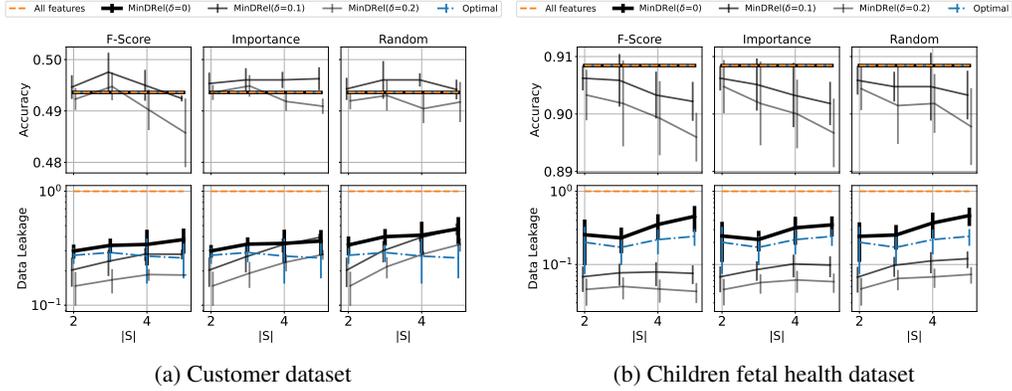


Figure 12: Comparison between using our proposed F-Score (left) with Importance (Middle) and Random (Right) for different choices of number of sensitive features $|S|$. The baseline classifier is a multinomial Logistic Regression

750 **Nonlinear classifiers** Similarly, we present a comparison of our proposed algorithms with the
 751 baseline methods when using non-linear classifiers. These metrics are reported for the Customer and
 752 Children Fetal Health datasets in in Figures 13a and 13b, respectively. The results show that using
 753 MinDRel with a value of $\delta = 0$ results in a minimal decrease in accuracy, but significantly reduces
 754 the amount of data leakage compared to the Baseline method.

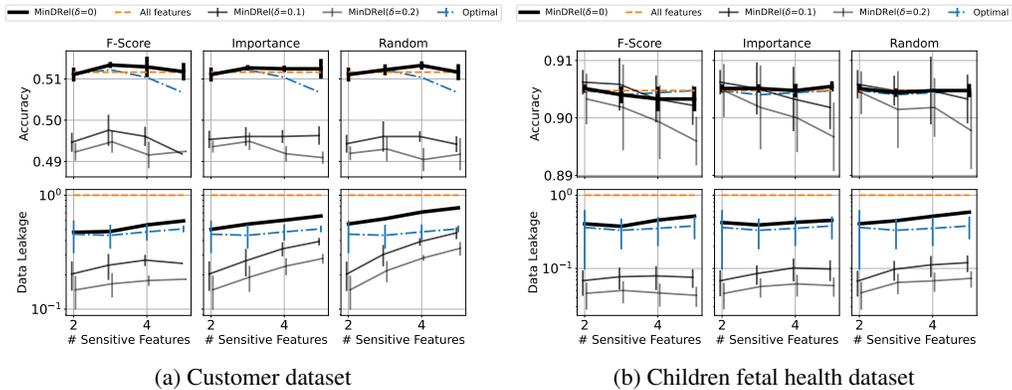


Figure 13: Comparison between using our proposed F-Score (left) with Importance (Middle) and Random (Right) for different choices of number of sensitive features $|S|$. The baseline classifier is a neural network classifier.