Winter 12-13-2018

The misty crystal ball: Efficient concealment of privacy-sensitive attributes in predictive analytics

Nicolas Banholzer
ETH Zurich

Stefan Feuerriegel
ETH Zurich

The misty crystal ball:
Efficient concealment of privacy-sensitive attributes in predictive analytics

Nicolas Banholzer[1]
ETH Zurich,
Zurich, Switzerland

Stefan Feuerriegel
ETH Zurich,
Zurich, Switzerland

ABSTRACT

Individuals are becoming increasingly concerned with privacy. This curtails their willingness to share sensitive attributes like age, gender or personal preferences; yet firms largely rely upon customer data in any type of predictive analytics. Hence, organizations are confronted with a dilemma in which they need to make a tradeoff between a sparse use of data and the utility from better predictive analytics. This paper proposes a masking mechanism that obscures sensitive attributes while maintaining a large degree of predictive power. More precisely, we efficiently identify data partitions that are best suited for (i) shuffling, (ii) swapping and, as a form of randomization, (iii) perturbing attributes by conditional replacement. By operating on data partitions that are derived from a predictive algorithm, we achieve the objective of masking privacy-sensitive attributes with marginal downsides for predictive modeling. The resulting trade-off between masking and predictive utility is empirically evaluated in the context of customer churn where, for instance, a stratified shuffling of attribute values impedes predictive accuracy rarely by more than a percentage point. Our proposed framework entails direct managerial implications as a growing share of firms adopts predictive analytics and thus requires mechanisms that better adhere to user demands for information privacy.

Keywords: information privacy; data masking; perturbation; predictive analytics; tree methods

INTRODUCTION

Personal information is valuable to online services (Bansal et al., 2010); however, individuals are increasingly concerned about their privacy. In fact, 82 percent of online users are reluctant to share personal information and more than a third has already given false information about themselves (Teltzrow and Kobsa, 2004). Information systems (IS) research has early noticed the importance of privacy for users (Hann et al., 2007) and the resulting demands by

[1] Corresponding author. nbanholzer@ethz.ch +41 44 632 84 64

practitioners for privacy-preserving technologies (e. g., Li and Raghunathan, 2014; Li and Sarkar, 2006, 2013).

Earlier research on privacy-preserving mechanisms can be distinguished by different goals, regarding which type of data is concealed (Duncan and Lambert, 1989). On the one hand, a set of methods has been developed in order to prevent identity disclosure, which is an important issue when data is supposed to be shared anonymously between a data owner and a third party (e. g., Li and Sarkar, 2013; Zhu et al., 2009). On the other hand, there is growing interest in so-called information privacy, which deals with individual concerns about confidentiality of certain sensitive attributes (Hann et al., 2007). This stream of literature acknowledges that certain data need to be stored for operational or legal reasons. Examples could consist of online shopping where billing addresses are needed for shipments or names need to be stored for tax audits. However, for sensitive attributes, such as age or gender, users expect organizations to be prudent. As part of their concerns about information privacy, users prefer sensitive attributes to be masked instead of being stored with their original value.

Prior work has addressed information privacy, but the methods for data masking focus on retaining important statistical properties (Li and Sarkar, 2006; Muralidhar and Sarathy, 2006), thereby ignoring the effect on predictive analytics. That is, privacy mechanisms must conceal attributes deemed sensitive, but still afford the use of predictive insights. However, there is reason to believe that even modest privacy gains almost completely destroy any analytical utility (Brickell and Shmatikov, 2008). It appears as if both information privacy and predictive utility cannot be achieved simultaneously and should rather be assumed mutually exclusive. This motivates our research as follows.

RESEARCH OBJECTIVE: *Designing a predictive framework for masking privacy-sensitive attributes in a given dataset while maintaining predictive power.*

In order to address our research objective, we propose a mechanism for protecting information privacy in predictive analytics. Thereby, we solve the decision problem of identifying attribute values that are best suited for masking by mitigating the effects on predictive performance. We approach this problem by partitioning the data into groups. Within each subgroup, masking is either accomplished by shuffling or swapping values between observations, or by replacing values with noisy ones from intervals that are clever chosen but can potentially be large. Our masking mechanism maintains some of the statistical properties from the underlying

distribution such as summary statistics in the case of shuffling. More importantly though, it makes a sparse use of data and perturbs sensitive attributes with little interferences concerning a given prediction task.

Our framework is evaluated by adapting the notion of regret to data masking, thereby obtaining a rigorous metric for assessing the loss in predictive power. As such, it complements established measures for gains in information privacy. Taken together, they support data-driven decision making by displaying the trade-off between information privacy and predictive analytics.

We perform a series of computational experiments in order to study how our proposed mechanism competes in a real-world setting. For this reason, we draw upon the task of predicting churn, as privacy-sensitive attributes are known to be associated with considerable prognostic capacity. This links to a direct trade-off between information privacy and predictive utility. For instance, our swapping mechanism masks an unknown portion of approximately 80 % of all values from a sensitive attribute, yet the changes in prediction performance (i. e., area under the curve) remain below a single percentage point. We further compare our mechanism for data masking across a variety of base classifiers in order to demonstrate that it consistently accomplishes only marginal effects on our regret metric, which thus contributes to the generalizability of mechanisms.

Our proposed prediction framework entails direct implications for individuals, IS practitioners and managers. First, individuals benefit from the fact that sensitive attributes become concealed before the data is permanently stored and eventually analyzed. Second, our predictive framework facilitates firms that want to strengthen information privacy. Here it allows firms to strategically choose a desired balance between a scarce use of sensitive data and expected benefits from predictive analytics. Third, we find that information privacy and firm utility are not incompatible and, with the necessary care in formulating the underlying decision problem, both objectives can be achieved simultaneously. In fact, our regret metric enables practitioners to assess the effects of perturbing sensitive information on firm utility quantitatively.

The paper is structured as follows. Section 2 summarizes earlier attempts in privacy-preserving analytics, revealing that these works are mostly concerned with anonymity and the potential risk of re-identifying individuals, while there is a scarcity of mechanisms with the

ability of data masking subject to preserving predictive utility. As a remedy, Section 3 proposes our decision problems that reconciles both goals: predictive modeling with strategic data masking. These mechanisms are then evaluated using a real-world study where churn is predicted (Section 4). Our findings lends to direct implications for management and IS practice as discussed in Section 5, while Section 6 concludes.

BACKGROUND

Information Privacy

There are several ways for firms to address privacy concerns. In recent years, differential privacy has emerged as a rigorous solution to protect personal information by giving mathematical guarantees about disclosure risks (Dwork, 2008). However, such high standards leave aside the fact that privacy demands are heterogeneous with some individuals guarding their privacy and others seeking convenience in more personalized service offerings through sharing personal information (Hann et al., 2007). It is possible for organizations to address this heterogeneity by catering to an individual's perception of how sensitive data is being handled. For instance, information practices (Culnan and Armstrong, 1999) and user control (Fusilier and Hoyer, 1980; Stone et al., 1983) can reconcile different demands for information privacy. Accordingly, we define information privacy as perceivable organizational efforts to conceal sensitive data about individuals, irrespective of whether firms themselves or the user controls data dissemination.

By definition, our work concerns the relationship between organizations and individuals from which they collect data. Individuals – in the sense "indivisible" – are uniquely identified by some attributes as name, phone or social security number, the so-called identifying attributes. Though anonymity matters a great deal to individuals, what they are more worried about is that their sensitive attributes, such as salary, medical test rest results or sexual orientation, are not revealed and related to their identity (Li and Raghunathan, 2014). Accordingly, privacy literature distinguishes two kinds of disclosures: identity and attribute disclosure (Duncan and Lambert, 1989). Preventing identity disclosure is a considerable issue when releasing data for analytical purposes to secondary users such as researchers (Narayanan and Shmatikov, 2008).

In order to address this issue, a large stream of literature began developing models which minimize the re-identification risk when disseminating data (e. g., Li et al., 2007; Machanavajjhala et al., 2006; Sweeney, 2002). However, the data owner himself often needs to retain the identifiability of the individuals in his database. Therefore, keeping the risk of identity

disclosure in mind, privacy-aware firms are also searching for mechanisms to mitigate the disclosure of sensitive attributes.

Privacy-Preserving Data Masking

Adam and Worthmann (1989) classify masking algorithms in statistical databases into four categories: conceptual, query restriction, data perturbation and output perturbation. For analytics tasks, research has found that not all are applicable (Zhu et al., 2009). Whilst conceptual approaches matter when designing privacypreserving data mining algorithms, our focus is on data perturbation techniques. Nowadays, perturbation techniques are often categorized under a more general notion of data masking techniques that include techniques beyond perturbation (Muralidhar and Sarathy, 2006). Below we summarize some concepts of masking models and related privacy-preserving algorithms.

Carefully derived algorithms have been found to establish a very strong from of privacy. Yet even though efficient solutions exist, they are often impractical (Yang et al., 2005). Furthermore, cryptographic approaches rather deal with the problem of sharing or jointly analyzing distributed databases with external parties, whereas the focus of this paper is the direct relationship between a user and an organization.

Microaggregation follows the idea of aggregating micro data into broader groups that provide anonymity to the individual on the group level. Some key concepts of anonymization based on microaggregation include *k-anonymity* (Samarati and Sweeney, 1998), *l-diversity* (Machanavajjhala et al., 2006) and *t-closeness* (Li et al., 2007). Data is said to be *k*-anonymous if an individual record cannot be distinguished from at least $k - 1$ individual records. Unfortunately, if a group of *k* individual records are indistinguishable, then an adversary is able to ascertain confidential values of a known member even from a homogeneous group. Therefore, grouped data is said to be additionally *l*-diverse if – simply put – there are at least *l* distinct values of the confidential attribute in each group. However, if confidential attributes are subject to highly skewed distributions, *l*-diversity can result in groups where the distribution of the confidential attribute differs substantially from the overall distribution. In this case, analytical findings can turn out to be misleading. Taking this problem into account, *t*-closeness requires that the distance between the group and overall distribution of a confidential attribute does not exceed a certain threshold *t*.

LeFevre et al. (2006) presents algorithms that satisfy *k*-anonymity, optionally *l*-diversity, but remain useful for classification or regression tasks. However, Brickell and Shmatikov (2008) finds that their application is in fact destructive to utility. Li and Sarkar (2006) propose an effective clustering-based method for masking categorical data. However, their method suffers from a disadvantage pertaining to most other microaggregation methods, which create groups where attribute similarity is greater within than between groups. This carries the risk of confidential attribute disclosure. In response, Li and Sarkar (2013) build a minimum spanning tree technique that creates groups where non-confidential attributes are similar, but confidential attributes are well distributed.

Rather than aggregating data, perturbation techniques randomize the individual records, e. g., by adding noise or shuffling attribute values. Shuffling preserves the marginal distribution of the data, whereas it depends on the desired strength of privacy how adding noise will alter the distribution. Both techniques have in common that inserting randomness will deprive data of information. Therefore, a large stream of research has been dedicated to build algorithms and models which transform data to ensure privacy and, simultaneously, keep it suitable for analytical purposes. Aggarwal and Yu (2008) provide a comprehensive survey of privacy-preserving data mining techniques, which includes – but is not limited to – perturbation.

The majority of privacy-preserving data masking is concerned with anonymity. However, the presented techniques, such as perturbation, are adaptable to serve a different goal of obscuring sensitive data as well (see e. g., Li and Sarkar, 2006; Muralidhar and Sarathy, 2006). This is important in the organizational context where identifying attributes need to be stored for operational or legal reasons, and anonymity is thus no
longer attainable.

Research Gap

Privacy mechanisms are commonly developed with the premise of ensuring anonymity. Such differentially private mechanisms establish scientific guarantees about the re-identification risk of individuals (Cynthia Dwork, 2006; Dwork, 2008). While vigorous protection of identity is a noble objective, it is overambitious in cases where organizations require identifying attributes for operational purposes. None the less, they may feel obliged to react to individual concerns about information privacy. For that reason, the field of privacy-preserving data mining has contributed various algorithms to prevent disclosure of sensitive information (Aggarwal and Yu, 2008).

Corresponding algorithms mask sensitive attributes and preserve statistical properties of the data (Li and Sarkar, 2006; Muralidhar and Sarathy, 2006). However, prior work does not take into account predictive analytics, which is a considerable issue because transforming data can be destructive to its utility (Brickell and Shmatikov, 2008). As a remedy, we develop a privacy mechanism that masks attributes deemed sensitive but with marginal downsides for predictive power.

Within the field of privacy-preserving data mining, there is no consistent way of quantifying the effects of data masking on data utility. But lacking a measurement puts firms at risk of forfeiting predictive utility without reaping the benefits of improved privacy. In this paper, we adopt the mathematical concept of regret to quantify the loss in predictive utility from transforming sensitive attributes under privacy considerations. A small regret informs firms about idle privacy potential, whereas a large regret warns them about destructive effects on predictive utility. Hitherto considered mutually exclusive, a comparison of regret emphasizes that there is a trade-off between between privacy and predictive analytics.

METHOD DEVELOPMENT

Intuition

While privacy is important for users (Hann et al., 2007), predictive analytics still plays a vital role for business (Fayyad et al., 1966). A frequent task in this discipline is classification, i. e., predicting the class label for an observation from its attributes. Masking attributes most likely impedes predictive accuracy because changing the attribute values corresponds to a loss in information. Thus, organizational efforts to strengthen information privacy in the form of data masking appear to run opposite to objectives in predictive analytics.

In order to resolve the conflicting goals of privacy and predictive utility, we postulate the idea that certain subsets of the data (i. e., so-called partitions) can be identified where, on the one hand, data masking is applied to conceal sensitive attributes. However, on the other hand, the prediction label for the observations in the partition are fairly similar and, hence, the masking operation has little effect on the overall predictive power. The underlying operations for masking sensitive attributes can then follow common strategies from the literature, namely, (i) shuffling, (ii) swapping, or (iii) replacing values.

Mathematically, we formalize the intuition behind our approach as follows. Let S refer to the sensitive attributes prior and $\tilde{S}$ after applying the masking operation. Further, let X denote all

remaining attributes which are not sensitive but nevertheless reported. Then, the objective in prediction is to estimate a model f that models Y based on S and X, i. e.,

$$f : (X, S) \mapsto Y. \tag{1}$$

The idea is that we partition the data $(X,S) = D_1 \cup ... \cup D_n$ into several disjoint subsets D_i, $i = 1,...,n$ with $D_i \cap D_j = \emptyset$ for $i \neq j$. Within each partition D_i, we expect that a different classifier f' that is trained with $\tilde{S}$ (instead of S) will obtain a similar outcome to f, i. e.,

$$f(D_i) \approx f'(D_i \mid_{S \leftarrow \tilde{S}}) \qquad \text{with } f' : (X, \tilde{S}) \mapsto Y. \tag{2}$$

Hence, the original predictive ability is retained for each partition.

Data Masking

Table 1 illustrates the effects of several data masking techniques on the original data. For instance, shuffling attribute values between observations is a particularly attractive technique in data masking because it perturbs the data without transforming the marginal distribution of an attribute (Muralidhar and Sarathy, 2006).

We first consider optimizing shuffling and swapping as permutation techniques for data masking. Let $s = (s_1,...,s_d)$ be a vector of $d = |D_i|$ expressions for sensitive attribute S_k in partition D_i and let $\Pi(s)$ be a permutation of s with $\mathcal{S}_d$ being the set of all possible permutations. Intuitively, we would want the permuted attribute values to be very different from the original ones. We achieve this goal by choosing the permutation that minimizes an equally weighted mean of Kendall's τ rank and Pearson's ρ linear correlation such that the average correlation approximates zero.

In contrast to permuting the original attribute values, replacing perturbs data through conditional replacement by sampling from an arbitrary distribution. For simplicity, we assume a uniform distribution, where the upper and lower bound correspond to the partition boundaries.[2]

Tree-Based Partioning

We suggest estimating a decision tree to accomplish our goal of masking sensitive attributes while mitigating the effects on predictive utility. This is because decision trees are very powerful in both partitioning and predicting, which previous studies have demonstrated in the context of privacy (Estivill-Castro and Brankovic, 1999; Li and Sarkar, 2009; Lindell and Pinkas,

[2] Replacing is used only for numerical variables. For categorical variables, we pertain to permutation techniques.

2002; Vaidya and Clifton, 2005). For instance, Li and Sarkar (2009) note their ability to mine sensitive information. In response, they propose a tree pruning approach for protection from tree-based classification attacks. Estivill-Castro and Brankovic (1999) observe that a permutation of the class labels within a data partition will not alter the assigned majority by the tree, but create uncertainty about the true class label of an individual. Finally, Lindell and Pinkas (2002) and Vaidya and Clifton (2005) develop tree-based privacy-preserving algorithms for horizontally and vertically distributed databases respectively.

In this work, we make use of the decision tree model developed in (Therneau et al., 1997). Based on recursive partitioning, data is recursively split into subgroups until a stopping criterion terminates the process. The path this observation traverses in the tree then determines the class label for an observation. Similar observations traverse the same path and thus each leaf node in the tree forms a group of observations; the data partition.

Small partitions or sharp constraints impede data masking because there is limited potential to permute or replace attribute values. This possibly leads to insufficient levels of information privacy. Naturally, the number of observations in a partition and the number of constraints are linked to the size of the decision tree. Recursive partitioning offers several ways for controlling the size of the tree. For instance, one can impose a minimum number of observations before making another split (Therneau et al., 1997). More commonly, a larger tree is grown first and subsequently pruned by removing splits that increase model complexity disproportionately. Intended to prevent overfitting (Therneau et al., 1997), pruning can also be employed to enforce stronger data masking.

Regret Metric for Predictive Performance

The objective of our masking mechanism is to solve a decision problem for organizations with a limited privacy budget. More formally, let D be the organization's database. After masking sensitive attributes in D with technique T_j, $j \in \{1,...,m\}$, the data is transformed to $\tilde{D}_j$. Now, assume the firm has a classification task where it wants to accurately predict the binary binary class label Y and denote with $\hat{Y}$ the predicted probability for the positive label. The following procedure is routine in Machine Learning. First, split D and $\tilde{D}_j$ each into a training and test set. Second, fit a model separately on the training sets of the original and transformed data. Third, predict $\hat{Y}$ on the test sets with $\hat{Y}$ being the predictions for the original and $\hat{Y}_{\tilde{D}_j}$ being the

predictions for the transformed data. We evaluate these predictions with the area under the receiver operating characteristic curve (AUC), which is a function of $\hat{Y}$ and Y.

A transformation of D is a decision to add privacy, potentially at the cost of firm utility. Therefore, the transformation should minimize the regret of making this decision. Measuring firm utility as the performance in the corresponding classification task, regret R is defined as the difference in AUC obtained using the original and transformed data

$$R_j = AUC\left(\hat{Y}_D, Y_D\right) - AUC\left(\hat{Y}_{\tilde{D}_j}, Y_D\right). \tag{5}$$

Intuitively, it measures the loss in predictive performance by applying masking technique T_j. The optimal masking technique T_{j*} is then found via

$$j^* = \underset{j=1,\ldots m}{\arg\min}\, R_j. \tag{6}$$

Since our masking mechanism is designed to maintain the predictive performance of the underlying classifier, we evaluate it more rigorously on a set of models; random forest, support vector machine as well as lasso, ridge and logistic regression (see following Section 4).

COMPUTATIONAL EXPERIMENTS

Empirical Setup

In our empirical evaluation, we draw upon the task of churn prediction as this presents a common, yet challenging, undertaking in IS research. It is important for the firm to know which individuals will continue to purchase and who will churn. For the latter group, the firm may want to approach them proactively, e. g., by offering specific deals. Hence, the task of predicting churn is directly linked to a case where predictive utility can obtain substantial financial benefits, but where users might not want to reveal their true personal information.

We specifically obtained a proprietary dataset as it better reflects the characteristics of data in real-world settings, such as a wide set of predictors and class imbalances. For that latter reason, we later draw upon the area under curve metric, as it is not sensitive to imbalances in the outcome variable. One might argue that the number and volume of purchases might present also very sensitive attributes. Therefore, we exclude these attributes from our masking algorithms.

Table 1 lists the model variables from our dataset. For the purpose of our analysis, we treat age and gender as sensitive attributes. We observe that most individuals are male and in their forties to fifties. Some individuals apparently registered on the online platform but did not make any

purchase. The dataset is fairly imbalanced, as continuing use appears for less than a third of the customer base.

The dataset is processed as follows. All variables with a skewed distribution are subject to a log-transformation (e. g., number and volume of purchases). The dataset is further split into a training (n = 283, i. e., 71 %) and test set (n = 110, i. e., 21 %) at the turn of a year in order to rule out potential confounders from seasonal effects.

Table 1: Model Variables.

Predictor variables		**Outcome variable**
Sensitive variables	Non-sensitive variables	
Age	Number of purchases	Binary label
Gender	Volume of purchases	*...continuing use: 110*
	Days since last purchase	*...churn: 273*
	Days since registration	
	Private/organizational customer	

Results

The computational results for the regret are presented in Table 2. It reveals that our masking mechanism hardly affects predictive accuracy. Masking the sensitive attributes often even leads to a slight increase in the performance of the classification task (negative regret). In case of shuffling, the performance boost is even significant. Nevertheless, the average change is still below a percentage point with shuffling being the least regretful masking technique. This is particularly a promising result as shuffling is also the only method that does not alter marginal distributions and summary statistics about the sensitive attributes.

Table 2: Prediction results under data masking.

Classifier	**Shuffling**		**Swapping**		**Replacing**	
	Mean (SD)	CI	Mean (SD)	CI	Mean (SD)	CI
Rf	-1.95** (1.94)	[-3.33; -0.56]	-0.69 (1.46)	[-0.35; 1.74]	-1.22* (1.85)	[-2.54; 0.11]
SVM	-0.39** (0.90)	[-1.04; 0.26]	-0.28 (1.07)	[-1.05; 0.49]	-0.09* (1.11)	[-0.88; 0.71]
Lasso	-0.80** (0.97)	[-1.49; -0.11]	-0.67 (0.92)	[-1.33; -0.04]	-0.30* (0.84)	[-0.93; 0.30]
Ridge	-0.70** (0.84)	[-1.30; -0.11]	-0.54 (0.82)	[-1.12; 0.05]	-0.31* (0.69)	[-0.80; 0.19]
Logit	-0.91** (0.90)	[-1.55; -0.27]	-0.44 (0.69)	[-0.93; 0.06]	-0.37* (0.84)	[-0.98; 0.23]
Average	-0.95*** (1.26)	[-1.31; -0.59]	-0.25 (1.10)	[-0.56; 0.07]	-0.46*** (1.17)	[-0.79; -0.13]

*Notes: The table shows the mean regret for m=10 simulations with the standard deviation in brackets. A mean regret different from zero at the 10 %; 5 % and 1 % significance level is indicated by *; ** or *** respectively. The 95 % confidence interval assumes a t-distribution with v =9 (model aggregate: v =49) degrees of freedom. Regret is computed as the percentage point change in the area under the curve after masking sensitive data. In case of replacing, missing lower and upper bounds are respectively set to the minimum and maximum value of the corresponding attribute. Other parameters in subgrouping are left at their defaults. For the evaluation models, hyper-parameters are tuned using 10-fold repeated cross-validation.*

DISCUSSION

Contributions to IS Research

Extant research is dedicated to develop mechanisms and algorithms that transform personal data to achieve privacy for the individual (Aggarwal and Yu, 2008). There are two streams of literature that can be divided by the goal of anonymization (Agrawal and Aggarwal, 2001;

Garfinkel et al., 2007; Li and Sarkar, 2013; Yang et al., 2005) and confidentiality about sensitive attributes (Hann et al., 2007; Li and Sarkar, 2006; Muralidhar and Sarathy, 2006). However, few link privacy-preserving methods to predictive analytics, which is concerning given the empirical evidence that privacy can be extremely costly for data analytics (Brickell and Shmatikov, 2008). To the best of our knowledge, none of these works explicitly develops a custom predictive algorithm for information privacy. Prior privacy mechanisms for data masking aim at minimizing the changes to distributional properties of the data (Li and Sarkar, 2006; Muralidhar and Sarathy, 2006). Predictions – if used at all – merely represents a form of evaluation but not the prime objective. From an algorithmic perspective, our masking mechanisms are designed to add perturbations where it is has little interferences with values that are decisive for forming predictions. Thus, our work addresses an imminent research gap in the area of privacy-preserving analytics as our proposed mechanism achieves high gains in information privacy, while simultaneously mitigating effects on the predictive utility.

Limitations and Potential for Future Research

Our masking mechanism yields no mathematical guarantees of how well the privacy of an individual is protected in the case of data breaches or other occasions where the data is shared. In fact, this objective is addressed in the context of differential privacy (Dwork, 2008; McSherry and Mironov, 2009), while our research acknowledges that firm-user interactions cannot work without user identification for legal and operational reasons. Instead, our mission is to limit the use of individual data with regard to certain attributes deemed as sensitive.

Implications for Management and IS Practice

Information privacy has considerable consequences for firms that want to benefit from the recent wave of predictive analytics (Chen et al., 2012). Prominent examples from our field include predictive analytics with the goal of targeted marketing, personalizing customer service offerings or predicting customer churn. Despite potential benefits, firms now face a dilemma where privacy-related features have become an integral element of online services, yet where potential improvements to information privacy remain untapped. For instance, bolstering information privacy could proof advertising to be more effective (Tucker, 2014) and positively influence customer purchase behavior (Tsai et al., 2011).

Our empirical findings are encouraging for firms seeking to accommodate both data sparsity and predictive utility. In fact, our results demonstrate that both are not necessarily

mutually exclusive. Rather, firms can successfully obfuscate sensitive attributes by perturbing sensitive attributes without lowering the prediction performance of the corresponding task. This finding directly benefits managers that want to implement similar approaches in order for individuals to benefit from a reduced disclosure of sensitive attributes.

The findings become especially relevant as "freemium" remains a dominant business model in online settings. Even though individuals have been found to be willing to pay for increased information privacy Hann et al. (2007), the majority rather prefer to save money when being confronted with the trade-off between sharing personal information and receiving a lower price (Beresford et al., 2012). Our work especially pertains to firms operating in such settings, as it might be feasible to increase information privacy with marginal downsides for data-driven decision-making.

In our work, we deliberately introduce the trade-off between privacy-sensitive attributes and predictive utility as a managerial decision problem. This allows firms to vary the extent of perturbation or shuffling depending on the underlying prediction task and the desired loss in prediction accuracy. Hence, firms can choose to maximize masking in order to attract privacy-sensitive individuals. Alternatively, they can prefer to bolster firm utility by luring individuals with personalized product and services. By incorporating our proposed regret metric into managerial decision-making, it is further possible for firms to quantify the effects of stronger data masking. This is important because it is difficult for firms to offer individuals more privacy when they can hardly estimate the value of utilizing less personal information. As a consequence, we suggest that firms should base their decision about the appropriate privacy mechanism on the regret of the respective data transformation.

REFERENCES

Adam, N. R., and Worthmann, J. C. 1989. "Security-control methods for statistical databases: A comparative study," *ACM Computing Surveys* (21:4), pp. 515–556.

Aggarwal, C. C., and Yu, P. S. 2008. "A General Survey of Privacy-Preserving Data Mining Models and Algorithms," in *Privacy-Preserving Data Mining*, vol. 34 of *Advances in Database Systems*, A. K. Elmagarmid, A. P. Sheth, C. C. Aggarwal, and P. S. Yu, (eds.), Boston, MA: Springer US, pp. 11–52.

Agrawal, D., and Aggarwal, C. C. 2001. "On the design and quantification of privacy preserving data mining algorithms," in *Proceedings of the 20th ACM SIGMOD-SIGACT-SIGART symposium on Principles of database systems*, New York, USA: ACM Press, pp. 247–255.

Bansal, G., Zahedi, F. M., and Gefen, D. 2010. "The impact of personal dispositions on information sensitivity, privacy concern and trust in disclosing health information online," *Decision Support Systems* (49:2), pp. 138–150.

Beresford, A. R., Kübler, D., and Preibusch, S. 2012. "Unwillingness to pay for privacy: A field experiment," *Economics Letters* (117:1), pp. 25–27.

Brickell, J., and Shmatikov, V. 2008. "The cost of privacy," in *Proceeding of the 14th ACM SIGKDD International Conference on Knowledge discovery and data mining*, New York, USA: ACM Press, pp. 70–78.

Chen, H., Chiang, R. H. L., and Storey, V. C. 2012. "Business Intelligence and Analytics: From Big Data to Big Impact," *MIS Quarterly* (36:4), pp. 1165–1188.

Culnan, M. J., and Armstrong, P. K. 1999. "Information Privacy Concerns, Procedural Fairness, and Impersonal Trust: An Empirical Investigation," *Organization Science* (10:1), pp. 104–115.

Cynthia Dwork 2006. "Differential Privacy," .

Duncan, G., and Lambert, D. 1989. "The Risk of Disclosure for Microdata," *Journal of Business & Economic Statistics* (7:2), pp. 207–217.

Dwork, C. 2008. "An Ad Omnia Approach to Defining and Achieving Private Data Analysis," in *Privacy, Security, and Trust in KDD*, vol. 4890 of *Lecture Notes in Computer Science*, F. Bonchi, E. Ferrari, B. Malin, and Y. Saygin, (eds.), Berlin, Heidelberg: Springer, pp. 1–13.

Estivill-Castro, V., and Brankovic, L. 1999. "Data Swapping: Balancing Privacy against Precision in Mining for Logic Rules," in *Data Warehousing and Knowledge Discovery*, vol. 1676 of *Lecture Notes in Computer Science*, G. Goos, J. Hartmanis, J. van Leeuwen, M. Mohania, and A. M. Tjoa, (eds.), Berlin, Heidelberg: Springer Berlin Heidelberg, pp. 389–398.

Fayyad, U., Piatetsky-Shapiro, G., and Smyth, P. 1966. "From data mining to knowledge discovery in databases," *AI magazine* (17:3), pp. 37–54.

Fusilier, M. R., and Hoyer, W. D. 1980. "Variables affecting perceptions of invasion of privacy in a personnel selection situation," *Journal of Applied Psychology* (65:5), pp. 623–626.

Garfinkel, R., Gopal, R., and Thompson, S. 2007. "Releasing Individually Identifiable Microdata with Privacy Protection Against Stochastic Threat: An Application to Health Information," *Information Systems Research* (18:1), pp. 23–41.

Hann, I.-H., Hui, K.-L., Lee, S.-Y. T., and Png, I. P. 2007. "Overcoming Online Information Privacy Concerns: An Information-Processing Theory Approach," *Journal of Management Information Systems* (24:2), pp. 13–42.

LeFevre, K., DeWitt, D. J., and Ramakrishnan, R. 2006. "Workload-aware anonymization," in *Proceedings of the 12th ACM SIGKDD International Conference on Knowledge discovery and data mining*, New York, USA: ACM Press, pp. 277–286.

Li, N., Li, T., and Venkatasubramanian, S. 2007. "t-Closeness: Privacy Beyond k-Anonymity and l-Diversity," in *Proceedings of the 23rd IEEE International Conference on Data Engineering*, IEEE, pp. 106–115.

Li, X.-B., and Raghunathan, S. 2014. "Pricing and disseminating customer data with privacy awareness," *Decision Support Systems* (59), pp. 63–73.

Li, X.-B., and Sarkar, S. 2006. "Privacy Protection in Data Mining: A Perturbation Approach for Categorical Data," *Information Systems Research* (17:3), pp. 254–270.

Li, X.-B., and Sarkar, S. 2009. "Against Classification Attacks: A Decision Tree Pruning Approach to Privacy Protection in Data Mining," *Operations Research* (57:6), pp. 1496–1509.

Li, X.-B., and Sarkar, S. 2013. "Class Restricted Clustering and Micro-Perturbation for Data Privacy," *Management Science* (59:4), pp. 796–812.

Lindell, and Pinkas 2002. "Privacy Preserving Data Mining," *Journal of Cryptology* (15:3), pp. 177–206.

Machanavajjhala, A., Gehrke, J., Kifer, D., and Venkitasubramaniam, M. 2006. "L-diversity: Privacy beyond k-anonymity," in *22nd International Conference on Data Engineering*, IEEE, p. 24.

McSherry, F., and Mironov, I. 2009. "Differentially private recommender systems: Building privacy into the Netflix Prize contenders," in *Proceedings of the 15th ACM SIGKDD international conference on Knowledge discovery and data mining*, New York, USA: ACM Press, p. 627.

Muralidhar, K., and Sarathy, R. 2006. "Data Shuffling: A New Masking Approach for Numerical Data," *Management Science* (52:5), pp. 658–670.

Narayanan, A., and Shmatikov, V. 2008. "Robust De-anonymization of Large Sparse Datasets," in *Proceedings of the IEEE Symposium on Security and Privacy*, IEEE, pp. 111–125.

Samarati, P., and Sweeney, L. 1998. "Protecting privacy when disclosing information: k-anonymity and its enforcement through generalization and suppression," .

Stone, E. F., Gueutal, H. G., Gardner, D. G., and McClure, S. 1983. "A field experiment comparing information privacy values, beliefs, and attitudes across several types of organizations," *Journal of Applied Psychology* (68:3), pp. 459–468.

Sweeney, L. 2002. "k-anonymity: A model for protecting privacy," *International Journal of Uncertainty, Fuzziness and Knowledge-Based Systems* (10:05), pp. 557–570.

Teltzrow, M., and Kobsa, A. 2004. "Impacts of user privacy preferences on personalized systems: A comparative study," in *Designing personalized user experiences in eCommerce*, pp. 315–332.

Therneau, T. M., Atkinson, and Elizabeth J. 1997. "An introduction to recursive partitioning using the RPART routines".

Tsai, J. Y., Egelman, S., Cranor, L., and Acquisti, A. 2011. "The Effect of Online Privacy Information on Purchasing Behavior: An Experimental Study," *Information Systems Research* (22:2), pp. 254–268.

Tucker, C. E. 2014. "Social Networks, Personalized Advertising, and Privacy Controls," *Journal of Marketing Research* (51:5), pp. 546–562.

Vaidya, J., and Clifton, C. 2005. "Privacy-Preserving Decision Trees over Vertically Partitioned Data," in *Data and Applications Security XIX*S. Jajodia and D. Wijesekera, (eds.), Berlin, Heidelberg: Springer Berlin Heidelberg, pp. 139–152.

Yang, Z., Zhong, S., and Wright, R. N. 2005. "Privacy-Preserving Classification of Customer Data without Loss of Accuracy," in *Proceedings of the SIAM International Conference on Data Mining*, Philadelphia, PA: Society for Industrial and Applied Mathematics, pp. 92–102.

Zhu, D., Li, X.-B., and Wu, S. 2009. "Identity disclosure protection: A data reconstruction approach for privacy-preserving data mining," *Decision Support Systems* (48:1), pp. 133–140.