DO NOT DISTURB? Classifier Behavior on Perturbed Datasets

Bernd Malle, Peter Kieseberg, Andreas Holzinger

► To cite this version:

Bernd Malle, Peter Kieseberg, Andreas Holzinger. DO NOT DISTURB? Classifier Behavior on Perturbed Datasets. 1st International Cross-Domain Conference for Machine Learning and Knowledge Extraction (CD-MAKE), Aug 2017, Reggio, Italy. pp.155-173, 10.1007/978-3-319-66808-6_11 . hal-01677128

HAL Id: hal-01677128

<https://hal.inria.fr/hal-01677128>

Submitted on 8 Jan 2018

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

Distributed under a Creative Commons Attribution| 4.0 International License

DO NOT DISTURB ?

Classifier behavior on perturbed datasets

Bernd Malle^{1,2}, Peter Kieseberg^{1,2} and Andreas Holzinger¹

¹ Holzinger Group HCI-KDD
Institute for Medical Informatics, Statistics & Documentation
Medical University Graz, Austria
b.malle@hci-kdd.org

² SBA Research gGmbH, Favoritenstrae 16, 1040 Wien
PKieseberg@sba-research.org

Abstract. Exponential trends in data generation are presenting today’s organizations, economies and governments with challenges never encountered before, especially in the field of privacy and data security. One crucial trade-off regulators are facing regards the simultaneous need for publishing personal information for the sake of statistical analysis and Machine Learning in order to increase quality levels in areas like medical services, while at the same time protecting the identity of individuals. A key European measure will be the introduction of the General Data Protection Regulation (GDPR) in 2018, giving customers the ‘right to be forgotten’, i.e. having their data deleted on request. As this could lead to a competitive disadvantage for European companies, it is important to understand which effects deletion of significant data points has on the performance of ML techniques. In a previous paper we introduced a series of experiments applying different algorithms to a binary classification problem under anonymization as well as perturbation. In this paper we extend those experiments by multi-class classification and introduce outlier-removal as an additional scenario. While the results of our previous work were mostly in-line with our expectations, our current experiments revealed unexpected behavior over a range of different scenarios. A surprising conclusion of those experiments is the fact that classification on an anonymized dataset with outliers removed in beforehand can almost compete with classification on the original, un-anonymized dataset. This could soon lead to competitive Machine Learning pipelines on anonymized datasets for real-world usage in the marketplace.

Keywords: Machine learning, knowledge bases, right to be forgotten, perturbation, k-anonymity, SaNGreeA, information loss, cost weighing vector, multi-class classification, outlier analysis, variance-sensitive analysis

1 Introduction and Related Work

In today’s data-driven industries which increasingly form the backbone of the 21st century’s economy, personal information is no longer only stored by private

companies, public service organizations or health providers. They also constitute a vital building-block for business intelligence and as a decision-making basis for improving services or public investments in measures for disease or natural disaster prevention. Therefore lies a crucial advantage in the publication, linkage, and systematic analysis of data sets from heterogeneous sources via statistics as well as Machine Learning. Any kind of institution which fails or is forbidden to engage in such activities, will in time face serious disadvantages on the marketplace or a lack in service quality compared to entities able to do so.

One specific challenge for data processing entities is increasingly imposed on them by the law. Under the new European General Data Protection Regulations (*GDPR*) taking effect on June 1st, 2018, customers are given a *right-to-be-forgotten*, meaning that an organization is obligated to remove a customer’s personal data upon request. For many organizations, this would incur serious additional investments and costs from their IT infrastructure, as even backup- or statistical systems must be connected, lest no ‘forgotten’ data will reappear. Nevertheless, the law will allow data analysis on anonymized datasets (for which a right-to-be-forgotten makes no sense from a technical point of view), so that organizations will soon be faced with the question: Do we learn on original data & bear all costs of the impeding bureaucracy, or shall we analyze anonymized datasets and risk significantly lower insights.

This brings us to the field of Privacy aware machine learning (PAML) [6], enabled and fostered by concepts like *k-anonymity* [20], in which a record is released only if it is indistinguishable from at least $k - 1$ other entities in the dataset. However, due to many personal records being high-dimensional in nature and *k-anonymity* being highly dependent on spatial locality (density) in order to effectively implement the technique in a statistically robust way, it might be difficult to anonymize data without suffering an intolerable amount of information loss [1]. Moreover, automatic dimensionality reduction might be helpful to preserve variance, but extracting the meaning, and therefore relevance, of arbitrary features would assist in making sense of the data with respect to a specific application domain [10].

Moreover, the original privacy requirement of *k-anonymity* [22] has over time been refined by concepts like *l-diversity* [16] (in which every equivalence group must contain at least l diverse sensitive values from the original dataset), *t-closeness* [15] (which prescribes that the local distribution over sensitive values within an equivalence group must not differ from its global distribution by more than a threshold t) as well as *delta-presence* [19] (which links the quality of anonymization to the risk posed by inadequate anonymization). Additionally, there is a whole discipline of measures summarized as *differential privacy* [7], which deals with methods of securely releasing sensitive information upon database queries by injecting controlled noise into responses.

As far as PAML is concerned, a comparison of different Machine Learning algorithms on anonymized datasets was already conducted in 2014 [24] by applying 6 different algorithms on 3 datasets, with very diverse results per algorithm. The main weakness of this paper is its usage of extremely differently-sized datasets

which does not easily allow comparison; moreover they only used one very low privacy setting of $k = 2$, preventing the authors from examining more interesting behavior as information content degrades further; this is a main point of our work.

The authors of [17] propose a scheme for controlling over-generalization of less identity-vulnerable QIs in diverse classes by determining the importance of QIs via Random Forest pre-computations as well as computing sensitive attribute diversity via the Simpson index [21]. Their resulting adaptive anonymization algorithm was compared to Mondrian [13] as well as IACk [14] and shows improvements w.r.t information loss as well as coverage (the number of descendant leaf nodes of generalized values in the taxonomy). Accuracy measured on classification tree, random forest and SVM shows equal or better performance when applied to a dataset anonymized by their proposed solution; it is interesting to note that their performance on large factors of k not only remains stable, but in some cases increases with k , the same behavior we also observed in some of our experiments.

A recent paper [12] proposes the introduction of an additional requirement for anonymization on top of k -anonymity called h -ceiling, which simply restricts generalizations within an equivalence class to a certain level below suppression. In the case on an equivalence class being able to satisfy h -ceiling but not k -anonymity (their method applies full-domain generalization), counterfeit records are inserted into the respective group; each insertion is also collected in a journal which is eventually published with the anonymized data. Their approach unsurprisingly yields lower reconstruction error and information loss as well as more fine-grained query results due to less generalization. However, their experiments mostly fix $k = 5$ and therefore simply try to reduce information loss due to anonymization, but do not try to examine ML performance over a wider range of k factors; moreover, there seems to be some inconsistency in their predictions.

Finally, we should also reference our previous work on this topic [18], in which we conducted a comparison study of binary classification performance on perturbed (selective deletion) vs. wholesale anonymized data. Our experiments showed that perturbation was still significantly less damaging to Machine Learning performance than even slight anonymization; that state of our previous research marks the connecting point to this paper.

2 K-Anonymity and Information loss

While there are several data-structures which can contain and convey personal information we might want to protect (free text, audio, images, graph structures etc.) we are focusing our work on tabular data, since most unstructured documents of sensitive nature today can be mapped to tabular data and since delicate information is most easily extracted from those. Figure 1 illustrates the original tabular concept of three different categories of data we will encounter in such tables:

- **Identifiers** directly reveal the identity of a person without having further analysis of the data. Examples are first and last names, email address or social security number (SSN). As personal identifiers are hard to generalized (see Figure 3) in a meaningful way (truncating an email address to 'host' would not yield much usable information), those columns are usually removed. The figure displays this column in a red background color.
- **Sensitive data**, or 'payload', is crucial information for statisticians or researchers and can therefore not be erased or perturbed; such data usually remains untarnished within the released dataset. The table shows one column in green background color representing such data.
- **Quasi identifiers (QI's)**, colored in the table with an orange background, do not directly identify a person (age=35), but can be used in combination to restrict possibilities to such a degree that a specific identity follows logically. For instance, [23] mentioned that 87% of U.S. citizens in 2002 could be re-identified by just using the 3 attributes *zip code*, *gender* and *date of birth*. On the other hand, this information might hold significant information for the purpose of research (e.g. zip code could be of high value in a study on disease spread). Therefore we generalize this kind of information, which means to lower its level of granularity. As an example, one could generalize grades from A+ to B- into A's and B's and then further up to encompass 'all' (also denoted as '*'), as shown in Figure 3.

Name	Age	Zip	Gender	Disease
Alex	25	41076	Male	Allergies
...	...	...	...	...

Fig. 1. The three types of data considered in (k-)anonymization

As described in [5], k-anonymization requires a data release to contain at least $k - 1$ duplicate entries for every occurring combination of attributes. One can imagine this as a clustering problem with each cluster's (also called *equivalence class*) quasi-identifier state being identical for every data point it contains. One can achieve this via suppression and generalization, where by suppression we mean simple deletion, whereas generalization refers to a decrease in a value's granularity. As an example, in Figure 2, an input dataset has been transformed through k-anonymization into a clustered set with each cluster being at least of *size* = 3; thus the data is said to be 3 – *anonymized*.

Generalization works through a concept called *generalization hierarchies / taxonomies*, which run from leaf nodes denoting particular values ('France') via internal nodes ('Western Europe') to their most general root ('all countries' or '*'). Such a hierarchy is depicted in Figure 3. In generalizing the original input

Node	Name	Age	Zip	Gender	Disease
X1	Alex	25	41076	Male	Allergies
X2	Bob	25	41075	Male	Allergies
X3	Charlie	27	41076	Male	Allergies
X4	Dave	32	41099	Male	Diabetes
X5	Eva	27	41074	Female	Flu
X6	Dana	36	41099	Female	Gastritis
X7	George	30	41099	Male	Brain Tumor
X8	Lucas	28	41099	Male	Lung Cancer
X9	Laura	33	41075	Female	Alzheimer

Node	Age	Zip	Gender	Disease
X1	25-27	4107*	Male	Allergies
X2	25-27	4107*	Male	Allergies
X3	25-27	4107*	Male	Allergies
X4	30-36	41099	*	Diabetes
X5	27-33	410**	*	Flu
X6	30-36	41099	*	Gastritis
X7	30-36	41099	*	Brain Tumor
X8	27-33	410**	*	Lung Cancer
X9	27-33	410**	*	Alzheimer

Fig. 2. Tabular anonymization: input table and anonymization result

value, one traverses the tree from a leaf node upwards until a certain condition is met. In the case of k -anonymity, we satisfy this condition when we can construct an equivalence group with all quasi-identifiers being duplicates of one another.

Figure 1: A possible generalization hierarchy for the attribute “Quality”.

Fig. 3. Example of a typical generalization hierarchy taken from [2]

As each level of generalization invokes an increasing loss of specificity, we do not want to construct our clusters inefficiently, but minimize a dataset’s overall information loss [2]. This makes k -anonymization an NP-hard problem due to an exponential number of possible data-row combinations one can examine.

3 Experiments

The following sections will describe our series of experiments in detail, encompassing the dataset used, the algorithms chosen for classification as well as a description of the overall process employed to obtain our results.

3.1 Data

As input data we chose the training set of the adults dataset from the UCI Machine Learning repository which was generated from US census data and contains approximately 32,000 entries (30162 after deleting rows with incomplete information). All but one columns were considered for experimentation, the remaining representing duplicate information (education => education_num). Figure 4 shows the attribute value distribution of 6 arbitrarily selected columns of the original (un-anonymized) dataset.

Fig. 4. Initial distribution of six selected data columns of the adult dataset.

Amongst these distribution, two clearly stand out: *native-country* as well as *hours-per-week*, which are both dominated by a single attribute value (*United-States* and *40*, respectively). In order to demonstrate the effect of anonymization on attribute value distributions, Figure 5 shows the same attribute distributions under anonymization by a factor of $k = 19$. Although the dominance of the

United-States was successfully ”broken” by this method, in several instances the *generalized-to-all-value* (*) now skews the data set even more. In addition to the incurred information loss this might be another reason for degraded classifier performance on such data.

Fig. 5. Anonymized distribution of six selected data columns of the adult dataset, anonymization factor $k = 19$, with equal weight for each attribute.

3.2 Anonymization Algorithm

We implemented our own version of a greedy clustering algorithm called SaN-GreeA (Social network greedy clustering, [4]) in JavaScript mainly for three reasons: 1) apart from ’normal’ tabular anonymization it has a network anonymization component based on stochastic reconstruction error, so it is possible for us to use this algorithm in later works regarding the impact of anonymization on graph algorithms; 2) we wanted a simple conceptual model so we could interact

with the algorithm and thus conduct interactive Machine Learning experiments in the future (those experiments are well under way at the time of this writing); 3) we wanted an algorithm capable of running in the browser so we could run our experiments online especially w.r.t. 2). The main downside of this choice is the reduced algorithmic performance of $O(n^2)$ as well as a further slow-down for JS vs native code of a factor of about 3 – 4. In the future, we will strive to implement faster algorithms which nevertheless retain properties suitable for our needs, narrowing down the simplicity - performance trade-off.

As mentioned, SaNGreeA consists of two strategies for tabular as well as network anonymization, with two respective metrics for information loss. The *Generalization Information Loss* or *GIL* consists of a categorical as well as a continuous part, with the former measuring the distance of a level-of-generalization from it's original leaf node in the generalization hierarchy (taxonomy), while the latter measures the range of a continuous-valued generalization (e.g. age cohort [35-40]) divided by the whole range of the respective attribute (e.g. overall age-range [17-90]).

$$\text{GIL}(cl) = |cl| \cdot \left(\sum_{j=1}^s \frac{\text{size}(\text{gen}(cl)[N_j])}{\text{size}(\min_{x \in N}(X[N_j]), \max_{x \in N}(X[N_j]))} + \sum_{j=1}^t \frac{\text{height}(\Lambda(\text{gen}(cl)[C_j]))}{\text{height}(H_{C_j})} \right)$$

where:

- $|cl|$ denotes the cluster cl 's cardinality;
- $\text{size}([i1, i2])$ is the size of the interval $[i1, i2]$, i.e., $(i2 - i1)$;
- $\Lambda(w), w \in H_{C_j}$ is the sub-hierarchy of H_{C_j} rooted in w ;
- $\text{height}(H_{C_j})$ denotes the height of the tree hierarchy H_{C_j} ;

The total generalization information loss is then given by:

$$\text{GIL}(G, S) = \sum_{j=1}^v \text{GIL}(cl_j)$$

And the normalized generalization information loss by:

$$\text{NGIL}(G, S) = \frac{\text{GIL}(G, S)}{n \cdot (s + t)}$$

As for the networking-part of this algorithm, it introduces a measure called *structural information loss* (SIL). The SIL is composed of two different components, which represent statistical errors of 1) intra-cluster as well as 2) inter-cluster reconstruction.

For the exact mathematical definitions of SIL & NSIL the reader is kindly referred to the original paper. Because the structural information loss cannot be computed exactly before the assembly of all clusters is completed, the exact

computations were replaced by the following distance measures:

Distance between two nodes:

$$\text{dist}(X^i, X^j) = \frac{|\{l | l = 1..n \wedge l \neq i, j; b_l^i \neq b_l^j\}|}{n - 2}$$

Distance between a node and a cluster:

$$\text{dist}(X, cl) = \frac{\sum_{X^j \in cl} \text{dist}(X, X^j)}{|cl|}$$

Since SaNGreeA follows the greedy-clustering paradigm, it runs in quadratic time w.r.t. the input size in number of nodes. This worked well within milliseconds for a problem size of a few hundred nodes, but took up to 60 minutes on the whole adult training dataset. Finally, as stated above, we chose SaNGreeA for its intuitive simplicity and graph anonymization capabilities, the latter of which are serving us well in a different branch of our ongoing research efforts; for the experiments in this paper, we restricted ourselves to the tabular anonymization capabilities of the algorithm.

3.3 Dataset creation

To examine the effect of perturbation, anonymization, outlier-removal as well as outlier-removal+anonymization on classifier performance, we designed the following processing pipeline:

1. Taking the original (preprocessed) dataset as input, we transformed its attributes to boolean values, so instead of *native-country - > United-States* we considered *United-States - > yes / no*.
2. We ran 4 different classifiers on the resulting data and computed their respective F1 score. The 4 classifiers used were *gradient boosting* representing the boosting paradigm, *random forest* representing the bagging technique, *logistic regression* as a representative of categorical prediction via optimization of a coefficient vector, as well as *linear SVC* representing Support Vector Machines constructing hyperplanes in sufficiently high-dimensional spaces.
3. For our perturbation experiments, we extracted the most / least significant attribute values according to the logit coefficients as depicted in Figure 6. For each of these attribute values, we subsequently deleted a specific percentage $p \in \{0.2, 0.4, 0.6, 0.8, 1.0\}$ of data rows containing that value, resulting in a series of new datasets of reduced size.
4. In order to measure the effects of k-anonymization on classifier performance, we applied SaNGreeA's GIL component to generate datasets with a k-factor of $k \in \{3, 7, 11, 15, 19, 23, 27, 31, 35, 100\}$. Furthermore, we used each of these settings with 3 different weight vectors: 1) equal weights for all attributes, 2) age information preferred ($\omega(\text{age}) = 0.88$, $\omega(\text{other_attributes}) = 0.01$) and 3) race information preferred ($\omega(\text{race}) = 0.88$, $\omega(\text{other_attributes}) = 0.01$).

5. The outlier-removal datasets were created by executing scikit-learn's Isolation-Forest in order to identify and remove outliers in a range of 5% – 95% (step-size: 5%) from the original dataset. This resulted in 18 new datasets for analysis.
6. Finally, in order to analyze classifier performance on an outlier/anonymization combination, we repeated the procedure described for anonymization on a surrogate dataset that had 30% of its outliers removed in beforehand.

Fig. 6. Attribute values within the adult dataset which contribute highest / lowest certainty to the classification of income (truncated at 1.0). The rightmost columns represent information which enable a classifier to discern most clearly between classes, while the leftmost columns (depending on their actual score) could even confuse the algorithm. We chose this example because income is a binary decision, so the values don't change per category to predict.

4 Results & Discussion

4.1 Perturbed Datasets - Selective Deletion

In order to be able to compare the impact of selectively deleting the most / least important attribute values (in fact, the whole data points containing those values) on different classifiers, we chose to select these values via examining the logit coefficients produced during logistic regression. Although this possibly entails non-erasure of the values specifically significant for each classifier, we chose algorithmic comparison as the more insightful criterion; the implicit assumption that the same attribute values would influence all classifiers approximately equally was largely confirmed by our results.

In contrast to binary classification, determining the 'right' values to delete for a multi-class problem is not always possible: Values contributing highly to the decision boundary for one class might be less significant in the case of another - accordingly one would expect inconclusive behavior in the case of a target for which the highest / lowest log coefficients do not line up over class boundaries.

For each of the targets 'marital-status' and 'education-num' we measured those interesting coefficients in the hope of improving / degrading algorithmic performance; that means deletion of highest logit's is supposed to remove certainty from an algorithm and decreasing performance, while deletion of lowest logit's is supposed to remove uncertainty, thus improving performance. Our analysis showed that while 'marital-status' had mainly the same most / least significant logit's across all classes, the attribute values for 'education-num' were rather diverse in this area.

In the latter case this lead to erratic behavior of the resulting performance curves, as can be seen in (Figure 7). It is interesting to note that 'income_ >50k' obviously held much larger significance for Logistic Regression than for the other classifiers, as their results showed f1 score improvement with this particular value eviscerating.

In the case of 'marital-status' almost the same attribute values were rated as most / least significant across all classes - this results in very clear outputs with the erasure of highly important values decreasing performance drastically while deletion of confusing values leading to a significant increase in classifier performance (Figure 8). While it is not surprising that relationship information shows high correlation with marital status, the opposite effects of *sex_Female* and *sex_Male* stand out as a slight curiosity - being a woman in this dataset seems to point less distinctly to a specific marital status than being a man.

4.2 Anonymized Datasets

Analogous to our previous work [18] we performed anonymization on the adult dataset for a range of values of k , but this time extending the range to $k \in \{3, 7, 11, 15, 19, 23, 27, 31, 35, 100\}$ for a broader observational basis of algorithmic behavior, especially towards higher values of k , as already conducted by other researchers [17], [12]. As we set out to examine multi-class classification

Fig. 7. Multi-class classification on target *education-num* under perturbation by selective deletion of the most / least contributing attribute values. Since different values are significant for deciding on different classes of education level, progressive deletion of this data results in indeterminate behavior.

performance, we chose the 'marital-status' and 'education-num' columns of the adult dataset as targets, treating income as an independent input feature. For 'marital-status' we left the 7 categorical values in the original dataset unchanged, whereas we clustered the 16 continuous 'education-num' levels into the 4 groups 'elementary school', 'high school including graduate', 'college up to Bachelors' as well as 'advanced studies'.

Our observation generally show the same type of behavior than in our previous experiments on target *income*, with one notable exception: The Random Forest classifier shows a sharp drop in algorithmic performance when operating on the very skewed 'age' and 'race' feature vectors, only to recover its discriminative power and increase in performance up to a k of 100. We also note a somewhat similar behavior for Logistic Regression, albeit not as distinctly. A possible explanation for this behavior could lie in the *bagging*-nature of Random Forest, meaning that the algorithm bootstraps by randomly sampling data-points from the overall population into possibly overlapping bags of 'local' data. As larger swaths of the input data become more and more equal with increasing levels of k , this would lead to less local over-fitting, thus making the job easier for a

Fig. 8. Multi-class classification on target *marital-status* under perturbation by selective deletion of the most / least contributing attribute values. Since the same values are significant for deciding different classes of marital status, progressive deletion leads to orderly increase / decrease of ML performance.

global averaging-strategy to filter out variance and improve generalization ability. However, if this was true, the maximum performance should not be recorded on the original (un-anonymized) dataset, thus we are currently at a loss of an adequate explanation for this specific case.

Classifier performance on target *marital-status* displayed the same basic behavior as above, including the mysterious conduct of the Random Forest in case of our age- and race-vectors. Moreover, the classification results are generally better than for *education-num*, which is probably caused by our somewhat arbitrary clustering of education levels during pre-processing. All in all, the pure anonymization-related results were almost in line with our expectations; in addition, our previous assessment that implementation of the ‘right-to-be-forgotten’ for individual users is preferable to wholesale anonymization, has not changed for the multi-class case.

4.3 “Outliers” removed

One question we didn’t tackle in our previous work was the one of outlier removal; this is relevant due to the fact that e.g. people showing abnormal behavior could

Fig. 9. Multi-class classification on target *education-num* on the adult dataset under several degrees of k-anonymization.

be supposed to exercise their 'right-to-be-forgotten' more frequently, especially in a social network scenario. For our experiments we chose the original adult dataset's income target, especially since we could thus directly compare the results with those of our previous work [18]. We used scikit-learn's Isolation-Forest classifier to identify outliers according to a given *contamination* level and performed an initial round of removing outliers in a range of 0.5% – 5%. Since ML performance decreased only marginally under those settings and we thus assumed that the dataset had been curated in such a way as to exclude significant outliers, we pivoted to a much broader investigation of examining classifier performance on a dataset with increasingly eviscerating variance. Thus we repeated the same procedure for "outlier" levels of 5% – 95%, gradually diminishing the dataset's size from over 30k to about 1.5k data points. In order to account for that dramatic reduction, we compared classifier behavior with a control instance of the adult dataset with the same levels of truncation, but under random deletion of data points, thus not targeting variance in the control set.

The results are shown in Figure 11 and exhibit similar behavior to the removal of most-significant attribute values in our previous work: While performance only decreases slightly for deletion levels under 55%, we see a dramatic drop over the

Fig. 10. Multi-class classification on target *marital-status* on the adult dataset under several degrees of k-anonymization.

second half of the range. The obvious explanation for this behavior lie in the fact that more homogeneous clusters of data make it harder for any algorithm to construct a decision boundary - though it is noteworthy that this applies to all 4 classifiers the same despite their fundamentally different approaches. Lastly, the comparison set shows no significant increase / decrease of performance over the whole range of data deletion, supporting our conclusion that decreasing data set size was not the dominating influence for the observed algorithmic behavior.

4.4 Anonymization on Outliers removed

One problem with outliers during anonymization is that it forces the algorithm to over-generalize attribute values; this can either happen towards the end-stages of a greedy-clustering procedure like SaNGreeA (in which case the damage might be limited to the outliers themselves), but could also influence a full-domain generalizing algorithm during determination of a whole column's suitable generalization level (in which case the whole dataset would suffer significantly higher information loss). This fact in combination with our previously described results based on outlier removal gave rise to an interesting possibility: what if we *combined* outlier removal with anonymization? On the one hand classifier perfor-

Fig. 11. Binary classification on target income based on a dataset with different degrees of outliers removed (= variance loss) vs. the same degree of data randomly deleted.

mance degrades with loss of variance, but for the very same reason information loss during anonymization might be limited to much more sufferable levels.

This led to our last round of experiments in which we took the adult dataset with 30% outliers removed and conducted k-anonymization as described in the respective earlier section (for time- and comparison reasons only on marital-status), the results of which can be seen in Figure 12. We were astonished to observe that - for the most part - classifiers performed better under this setting than under anonymization alone. For logistic regression, although age & race vectors performed worse then their anonymized-only counterparts, performance for equal weights was better for $k < 11$. With Random Forest, all vectors performed better than their anonymization-only counterparts, with $k = 3$ only 2% below original performance. With Linear SVC, age & race performed worse at the beginning only to recover with increasing performance towards $k = 100$, whereas the equal vector behaved about equal to it's non-outlier-removed opposite. Finally, Gradient Boosting in this setting outperforms it's anonymization-only competitor in all settings with it's $k = 3$ equal weight vector performance lying within only half a percentage point of the performance on the original, un-anonymized dataset.

As a side-note, we observe that under these settings, SVC starts to mimic Random Forest's behavior of an initial collapse in performance for the age- and range-vectors with a subsequent recovery towards higher levels of k . We do not yet have an adequate explanation for this and will investigate deeper in our future efforts.

Those amazing results raise a few burning questions: 1) Can we repeat that performance on real-world data? 2) Could we combine this technique with interactive Machine Learning / Anonymization which yield better weight vectors? 3) Do those advantages only hold for a toy algorithm or will they persist under more sophisticated Anonymization pipelines? 4) Can we further enhance those

Fig. 12. Multi-class classification on target marital status based on a dataset with 30% outliers removed AND under different degrees of k-anonymization.

results by mixing synthetic data into the dataset? 5) Will better feature engineering compensate for our original drop in performance and thus moot our insight? and 6) can we apply this conclusion to other data structures like social networks? These points shall now briefly be discussed before concluding the paper.

5 Open problems / Future challenges

1. **Real world data.** Despite the convenient availability of well-curated datasets with many thousands of data rows, actual industry datasets are usually orders of magnitudes larger. This has consequences for their internal data topology and thus the performance of ML algorithms; e.g. [3] observe that variance error can be expected to decrease as training set size increases (though this might have nothing to do with variance in the dataset itself).
2. **Interactive machine learning.** We have demonstrated experiments with different weight vectors in our approach regarding anonymization. However, data utility is highly subjective w.r.t. the specific area of application; therefore choosing the importance of attributes with regard to the particular environment is best done by a human. The problem of (k-)anonymization thus

lends itself to interactive Machine Learning (iML) with a human-in-the-loop approach [9], [11], [8]. We have implemented software for iML Anonymization and are currently collecting test results which will soon be ready for publication.

3. **Real world algorithms.** While we only anonymizing our datasets via simple k-anonymization through greedy clustering, there are much more sophisticated algorithms available, capable of fine-tuning generalization levels to the specific data topology of an input set in order to minimize information loss. It remains to be seen if such algorithms can still profit from removal of outliers as a pre-processing step.
4. **Synthetic datasets.** In recent years it has become common to augment (small) datasets via synthetically generated, additional data-points [25]. By controlling the data generation process, one would be able to also control variance-injection into a dataset. Therefore, instead of outlier-removal, one could enrich a dataset by introducing lower-variance data points before anonymization.
5. **Better feature engineering.** For our experiments, we considered practically all columns of the adult dataset, although some exhibited much higher variance than others. It is therefore conceivable that by careful feature engineering the basis for anonymization could be sufficiently improved, rendering outlier-removal unnecessary.
6. **Graph structure anonymization.** This includes questions of measuring structural outliers in a graph (maybe via centrality- or component-based analysis?) as well as outlier removal (do they have to be deleted or will randomly adding edges to such nodes suffice?). Our team is currently devising experiments in this direction, but our efforts are still in the early stages.

6 Conclusion

In this paper we continued our initial experiments on the effects of anonymization and perturbation of knowledge bases on classifier performance and expanded our efforts to multi-class classification, outlier-removal as well as a combined outlier/anonymization approach. Our results show that selective deletion of significant attribute values is preferable to general anonymization, insofar a dataset's topology allows for such conduct. We have furthermore seen that reducing variance in a dataset prevents algorithms of different breeds alike from finding efficient discriminators between classes, leading to a significant degradation of machine learning performance. Finally, we were astonished to observe that combining outlier-removal with anonymization can - under circumstances - yield almost as good a performance as classification on the original, un-anonymized dataset itself. We believe that this insight, in combination with work on interactive Anonymization we are currently conducting, state-of-the art anonymization techniques (we were using a rather simple algorithm for this paper), as well as the introduction of synthetic data, will enable us to soon propose competitive Machine Learning pipelines for real-world usage to counterbalance any regulatory disadvantage European companies are currently facing on the marketplace.

References

1. Charu C Aggarwal. On k-anonymity and the curse of dimensionality. In *Proceedings of the 31st international conference on Very large data bases VLDB*, pages 901–909, 2005.
2. Gagan Aggarwal, Tomas Feder, Krishnaram Kenthapadi, Rajeev Motwani, Rina Panigrahy, Dilys Thomas, and An Zhu. Approximation algorithms for k-anonymity. *Journal of Privacy Technology (JOPT)*, 2005.
3. Damien Brain and G Webb. On the effect of data set size on bias and variance in classification learning. In *Proceedings of the Fourth Australian Knowledge Acquisition Workshop, University of New South Wales*, pages 117–128, 1999.
4. Alina Campan and Traian Marius Truta. Data and structural k-anonymity in social networks. In *Privacy, Security, and Trust in KDD*, pages 33–54. Springer, 2009.
5. Valentina Ciriani, S De Capitani di Vimercati, Sara Foresti, and Pierangela Samarati. κ -anonymity. In *Secure data management in decentralized systems*, pages 323–353. Springer, 2007.
6. John C Duchi, Michael I Jordan, and Martin J Wainwright. Privacy aware learning. *Journal of the ACM (JACM)*, 61(6):38, 2014.
7. Cynthia Dwork. Differential privacy: A survey of results. In *International Conference on Theory and Applications of Models of Computation*, pages 1–19. Springer, 2008.
8. A Holzinger, M Plass, K Holzinger, GC Crisan, CM Pintea, and V Palade. Towards interactive machine learning (iml): Applying ant colony algorithms to solve the traveling salesman problem with the human-in-the-loop approach. In *IFIP International Cross Domain Conference and Workshop (CD-ARES)*, page in print. Springer, Heidelberg, Berlin, New York, 2016.
9. Andreas Holzinger. Interactive machine learning for health informatics: When do we need the human-in-the-loop? *Springer Brain Informatics (BRIN)*, 3(2):119–131, 2016.
10. Andreas Holzinger. Introduction to machine learning & knowledge extraction (make). *Machine Learning and Knowledge Extraction*, 1(1):1–20, 2017.
11. Peter Kieseberg, Bernd Malle, Peter Frhwirt, Edgar Weippl, and Andreas Holzinger. A tamper-proof audit and control system for the doctor in the loop. *Brain Informatics*, pages 1–11, 2016.
12. Hyukki Lee, Soohyung Kim, Jong Wook Kim, and Yon Dohn Chung. Utility-preserving anonymization for health data publishing. *BMC Medical Informatics and Decision Making*, 2017.
13. Kristen LeFevre, David J DeWitt, and Raghu Ramakrishnan. Mondrian multidimensional k-anonymity. In *Data Engineering, 2006. ICDE’06. Proceedings of the 22nd International Conference on*, pages 25–25. IEEE, 2006.
14. Jiuyong Li, Jixue Liu, Muzammil Baig, and Raymond Chi-Wing Wong. Information based data anonymization for classification utility. *Data & Knowledge Engineering*, 70(12):1030–1045, 2011.
15. Ninghui Li, Tiancheng Li, and Suresh Venkatasubramanian. t-closeness: Privacy beyond k-anonymity and l-diversity. In *IEEE 23rd International Conference on Data Engineering, ICDE 2007*, pages 106–115. IEEE, 2007.
16. Ashwin Machanavajjhala, Daniel Kifer, Johannes Gehrke, and Muthuramakrishnan Venkatasubramanian. l-diversity: Privacy beyond k-anonymity. *ACM Transactions on Knowledge Discovery from Data (TKDD)*, 1(1):1–52, 2007.

17. A Majeed, F Ullah, and S Lee. Vulnerability-and Diversity-Aware Anonymization of Personally Identifiable Information for Improving User Privacy and Utility of Publishing Data. *Sensors*, pages 1–23, 2017.
18. Bernd Malle, Peter Kieseberg, Edgar Weippl, and Andreas Holzinger. The right to be forgotten: towards machine learning on perturbed knowledge bases. In *International Conference on Availability, Reliability, and Security*, pages 251–266. Springer, 2016.
19. M. E. Nergiz and C. Clifton. delta-presence without complete world knowledge. *IEEE Transactions on Knowledge and Data Engineering*, 22(6):868–883, 2010.
20. Pierangela Samarati. Protecting respondents identities in microdata release. *IEEE Transactions on Knowledge and Data Engineering*, 13(6):1010–1027, 2001.
21. Edward H Simpson. Measurement of diversity. *Nature*, 1949.
22. Latanya Sweeney. Achieving k-anonymity privacy protection using generalization and suppression. *International Journal of Uncertainty, Fuzziness and Knowledge-Based Systems*, 10(5):571–588, 2002.
23. Latanya Sweeney. k-anonymity: A model for protecting privacy. *International Journal of Uncertainty, Fuzziness and Knowledge-Based Systems*, 10(05):557–570, 2002.
24. Hayden Wimmer and Loreen Powell. A Comparison of the Effects of K-Anonymity on Machine Learning Algorithms. pages 1–9, 2014.
25. Sebastien C Wong, Adam Gatt, Victor Stamatescu, and Mark D McDonnell. Understanding data augmentation for classification: when to warp? In *Digital Image Computing: Techniques and Applications (DICTA), 2016 International Conference on*, pages 1–6. IEEE, 2016.