Data Privacy Management
September 19, 2024

Balancing Privacy and Utility in Multivariate Time-Series Classification

Adrian-Silviu Roman, Béla Genge, and Piroska Haller

George Emil Palade University of Medicine, Pharmacy, Science, and Technology of Targu Mures
Romania

Outline

Introduction

- The collection and processing of sensor data involves important privacy issues.
- The case of Time-Series Classification (TSC).
- The objective is to protect the multivariate data by a combination of feature clustering and perturbation, and
- To balance the utility and privacy of the protected data.

Problem description

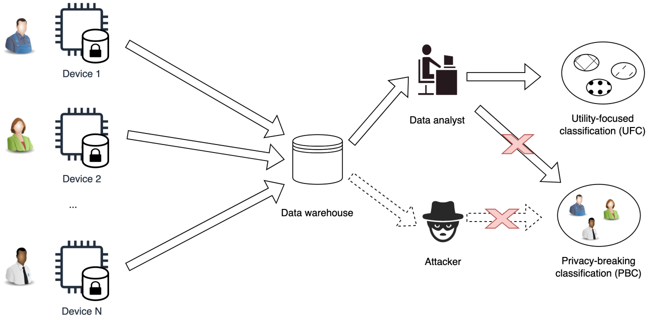

Figure 1: Problem description. Sensor data, collected from several devices, undergoes protection at the device level and is transmitted to a centralized data warehouse for classification, with utility and privacy constraints.

Problem description

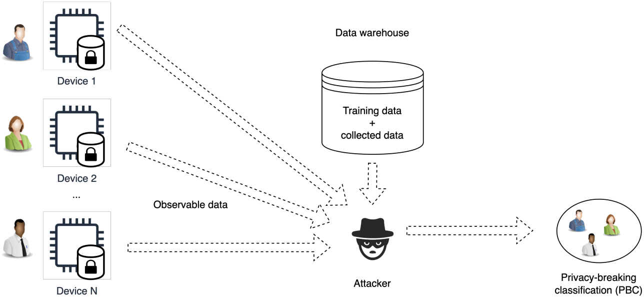

Figure 2: The adversarial model. The attacker has access to observable data, training data, and all previously collected data.

Problem description

Possible scenario:

- A fleet management company that collects data (vehicle speed, engine diagnostics, fuel consumption) from sensors in their vehicles.
- Data is sent to a third-party processor (an insurance company).
- Utility objectives:
  - **categorize the driving behavior**, distinguishing between aggressive and normal driving styles;
  - **categorize the type of road surfaces** used by the vehicles.
- Privacy objective:
  - **ensure that the data remains protected** from potential attackers or "honest-but-curious" actors (no user identification is possible).

Proposed approach

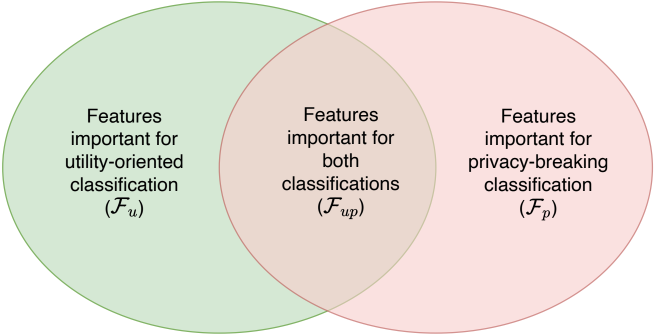

Figure 3: Feature clustering objective.

Proposed approach

- A dual-model, consisting of two competing classifiers, an **Utility-Focused Classifier** (UFC) and a **Privacy-Breaking Classifier** (PBC).
- A protection technique independent of the perturbation type, applying the perturbation to multivariate time-series data and utilizing feature importance to distribute the noise.
- The *classification utility-privacy balance* score, $\mathcal{B}_{UP}$.
- Experiments on two well-known driver datasets [1, 2] using the w-event LDP perturbation method.

Proposed approach

- Consider the classification models of UPC and PBC:

$$f_p : \mathbb{R}^{w \cdot d} \longrightarrow \mathcal{C}_u \text{ and } f_u : \mathbb{R}^{w \cdot d} \longrightarrow \mathcal{C}_p \tag{1}$$

- Define the classification accuracy on unperturbed data for each model as:

$$\mathcal{A}_u = \frac{1}{N \cdot T} \sum_{i=1}^{N} \sum_{t=1}^{T} \mathbb{I}(\mathbf{Y}_{ui}^t = f_u(\mathbf{X}_i^t)), \tag{2}$$

$$\mathcal{A}_p = \frac{1}{N \cdot T} \sum_{i=1}^{N} \sum_{t=1}^{T} \mathbb{I}(\mathbf{Y}_{pi}^t = f_p(\mathbf{X}_i^t)) \tag{3}$$

- Define the classification accuracy on perturbed data:

$$\mathcal{A}'_u(\theta) = \frac{1}{N \cdot T} \sum_{i=1}^{N} \sum_{t=1}^{T} \mathbb{I}(\mathbf{Y}_{ui}^t = f_u(\mathcal{M}(\mathbf{X}_i^t; \theta))), \text{ and} \tag{4}$$

$$\mathcal{A}'_p(\theta) = \frac{1}{N \cdot T} \sum_{i=1}^{N} \sum_{t=1}^{T} \mathbb{I}(\mathbf{Y}_{pi}^t = f_p(\mathcal{M}(\mathbf{X}_i^t; \theta))). \tag{5}$$

Proposed Approach

- The following conditions bound the data protection objective:

$$\mathcal{A}'_p(\theta) \ll \mathcal{A}_p \text{ and } \mathcal{A}'_u(\theta) \approx \mathcal{A}_u. \tag{6}$$

- We introduce $\mathcal{B}_{UP}$, the *classification utility-privacy balance*:

$$\mathcal{B}_{UP}(\theta) = 1 - \frac{\mathcal{A}'_u(\theta)}{\mathcal{A}_u} \cdot \left(1 - \frac{\mathcal{A}'_p(\theta)}{\mathcal{A}_p}\right). \tag{7}$$

- Optimisation objective - find the perturbation parameter set θ^* such that $\mathcal{B}_{UP}(\theta)$ is minimum:

$$\theta^* = argmin_\theta\{\mathcal{B}_{UP}(\theta)|\mathcal{B}_{UP}(\theta) > 0\}. \tag{8}$$

Proposed Approach

The proposed method for finding the perturbation parameters consists of the following steps:

1. Compute feature importance for the two classifications (UFC and PBC);
2. Cluster features based on the computed importance coefficients (using ρ_I):
$$\mathcal{F} = \mathcal{F}_u \cup \mathcal{F}_p \cup \mathcal{F}_{up} \tag{9}$$
3. Distribute and apply the perturbation (β_T) to the features in $\mathcal{F}_p$ and $\mathcal{F}_{up}$, using parameters α_p and α_{up};
4. Select the perturbation parameter set θ^* such that $\mathcal{B}_{UP}(\theta)$ is minimum, with $\theta = \{\rho_I, \beta_T, \alpha_p, \alpha_{up}\}$.

Proposition: Let $\mathcal{M}$ be a mechanism composed of m mechanisms $\mathcal{M}_1$, $\mathcal{M}_2$, ..., $\mathcal{M}_i$, ..., $\mathcal{M}_m$, $m < d$, one for each attribute/feature F_i ($F_i \in \mathcal{F}_p$ or $F_i \in \mathcal{F}_{up}$), satisfying ϵ_i-LDP, such that the privacy budget for each mechanism $\mathcal{M}_i$ is defined as follows:

$$\epsilon_i = \begin{cases} \frac{\alpha_p \cdot \beta_T}{|\mathcal{F}_p|}, & \text{if } F_i \in \mathcal{F}_p, \\ \frac{\alpha_{up} \cdot \beta_T}{|\mathcal{F}_{up}|}, & \text{if } F_i \in \mathcal{F}_{up}, \end{cases}$$

where $\alpha_p + \alpha_{up} = 1,$ with $\alpha_p, \alpha_{up} \in [0, 1].$ If the following condition is fulfilled:

$$\frac{\alpha_p}{|\mathcal{F}_p|} \leq \frac{\alpha_{up}}{|\mathcal{F}_{up}|}. \tag{10}$$

then the perturbation of features in $\mathcal{F}_p$ is higher than or equal to the perturbation applied on features in $\mathcal{F}_{up}$.

Experimental Results

Table 1: Classification accuracy for unprotected test data with a FCN-LSTM model.

Dataset	Classification objective	Achieved accuracy	Benchmark accuracy [3]
UAH Driveset [2]	Driver detection	0.9240	0.8986
	Behavior detection	0.8863	NA
	Road type detection	0.9998	NA
HCRL [1]	Driver detection	0.9120	0.9510
	Road type detection	0.9615	NA

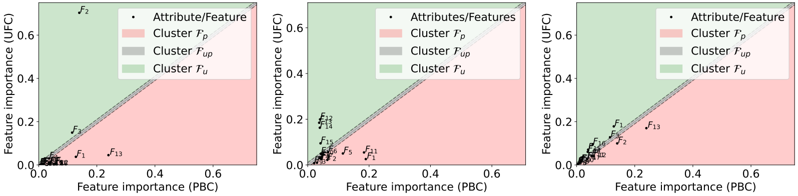

(a) Driver-vs-road-type (UAH Driveset)

(b) Driver-vs-road-type (HCRL Dataset)

(c) Driver-vs-behavior (UAH Driveset)

Figure 4: Feature clustering based on feature importance coefficients for two classifications (UFC and PBC), conducted using Random Forest with Gini importance ($\rho_I = 0.01$).

Experimental Results

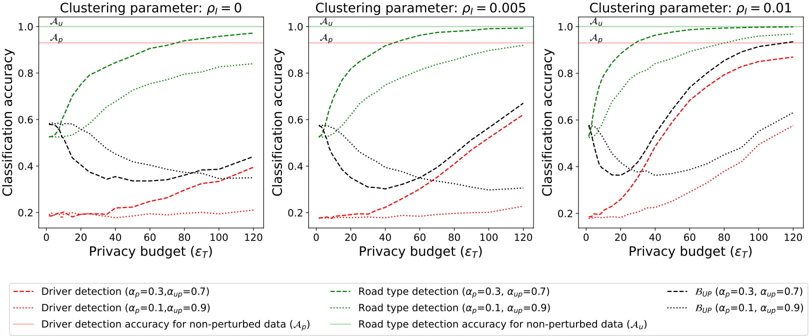

Figure 5: Driver-vs-road-type classification accuracy on perturbed data (UAH dataset).

Experimental Results

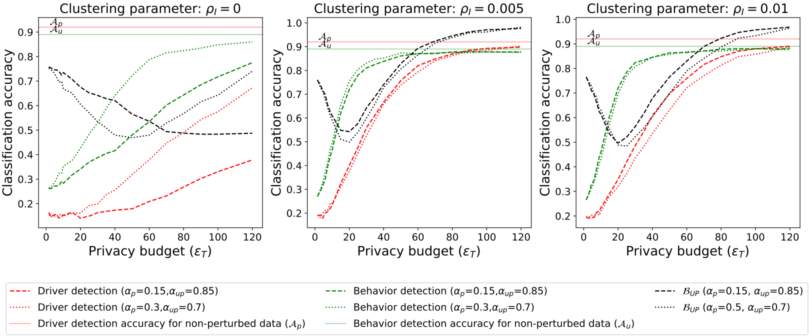

Figure 6: Driver-vs-behavior classification accuracy on perturbed data (UAH dataset).

Table 2: Classification accuracy for perturbed data using the proposed approach.

Dataset	Classification scenario	Perturbation approach	Clustering parameter (ρ_I)	# of features per cluster ($\mathcal{F}_p$, $\mathcal{F}_{up}$, $\mathcal{F}_u$)	Perturbation parameters (α_p, α_{up})	min($\mathcal{B}_{UP}$)	ϵ_T	$\mathcal{A}'_p(\theta)$	$\mathcal{A}'_u(\theta)$	MAE
UAH [2]	Driver-vs-road-type	w-event LDP (no clustering)	-	-	-	0.3348	15	0.2103	0.8610	1.0170
		w-event LDP (feature clustering, **proposed method**)	0	{15,0,2}	{0.3,0.7}	0.3357	60	0.2480	0.9058	0.7424
					{0.1,0.9}	0.3459	100	0.1941	0.8266	1.3350
			0.005	{12,3,2}	{0.3,0.7}	0.3022	40	0.2227	0.9175	0.6038
					{0.1,0.9}	**0.2973**	**100**	0.2012	0.8967	0.6921
			0.01	{10,5,2}	{0.3,0.7}	0.3630	20	0.2564	0.8831	0.8141
					{0.1,0.9}	0.3623	40	0.2253	0.8415	1.0194
HCRL [1]	Driver-vs-road-type	w-event LDP (no clustering)	-	-	-	0.5683	30	0.4356	0.7941	0.5291
		w-event LDP (feature clustering, **proposed method**)	0	{10,0,5}	{0.3,0.7}	0.5765	60	0.4613	0.8235	0.4306
					{0.15,0.85}	**0.5629**	**120**	0.4377	0.8076	0.4302
			0.01	{9,2,4}	{0.3,0.7}	0.5785	60	0.4635	0.8235	0.4304
					{0.15,0.85}	0.5749	100	0.3755	0.6945	0.5168
			0.015	{7,4,4}	{0.3,0.7}	0.5803	40	0.3841	0.6968	0.4946
					{0.15,0.85}	0.5707	90	0.4263	0.7750	0.4428
UAH [2]	Driver-vs-behavior	w-event LDP (no clustering)	-	-	-	0.4859	30	0.3090	0.6844	0.5077
		w-event LDP (feature clustering, **proposed method**)	0	{9,0,8}	**{0.3,0.7}**	**0.4692**	**50**	0.3181	0.7220	0.3651
					{0.15,0.85}	0.4834	100	0.3305	0.7175	0.3655
			0.005	{6,6,5}	{0.3,0.7}	0.4986	20	0.3759	0.7545	0.4110
					{0.15,0.85}	0.5426	20	0.4000	0.7201	0.5002
			0.01	{4,8,5}	{0.3,0.7}	0.4836	25	0.3714	0.7707	0.3609
					{0.15,0.85}	0.4965	20	0.3480	0.7207	0.4458

Conclusion

- We proposed a novel approach for protecting multivariate time series data in the context of TSC.
- The problem is defined in the context of two opposing classifiers (UFC and PBC).
- We introduced the *classification utility-privacy balance* score, $\mathcal{B}_{UP}$.
- The method achieves a balance between privacy preservation and data utility.

Balancing Privacy and Utility in Multivariate Time-Series Classification

Thank you!

Contact: adrian.roman@umfst.ro

References

Byung Il Kwak, Jiyoung Woo, and Huy Kang Kim.
Know your master: Driver profiling-based anti-theft method.
In *PST 2016*, 2016.

Eduardo Romera, Luis M Bergasa, and Roberto Arroyo.
Need data for driver behaviour analysis? presenting the public uah-driveset.
In *2016 IEEE 19th international conference on intelligent transportation systems (ITSC)*, pages 387–392. IEEE, 2016.

Abdellah El Mekki, Afaf Bouhoute, and Ismail Berrada.
Improving driver identification for the next-generation of in-vehicle software systems.
IEEE Transactions on Vehicular Technology, 68(8):7406–7415, 2019.