ON THE APPLICATION AND IMPACT OF ϵ-DP AND FAIRNESS IN AMBULANCE ENGAGEMENT TIME PREDICTION

Anonymous authors
Paper under double-blind review

ABSTRACT

This study elaborates on a complete pipeline for the development of a private and fair machine learning model to predict ambulance engagement time. Results indicated that sensitive variables reduced their impact on the creation of the model with RandomForestClassifier as the differential privacy budget (ϵ) was reduced. In addition, it was validated that by maintaining the same hyperparameter configuration and reducing the level of ϵ, the accuracy of the unprivileged group decreases until it stabilizes. Finally, it is possible to keep agents' and victims' privacy, recovering a machine learning model with good performance.

1 INTRODUCTION

In some countries, Fire Departments (FD) deal with a high demand for interventions and a shortage of resources Guyeux (2022). This results in long response times and leaves a locality unprotected. Motivated to optimize their resources, this study contributes to *i)* the development of a Machine Learning (ML) model to predict ambulance engagement time since rescue ambulances are their main resource, and to *ii)* the impact analysis of Differential Privacy (DP) Dwork (2006) and Reweighing (Rew) Calders et al. (2009) mechanisms on model performance, since there is sensitive information about victims and agents and because the training of the latter as volunteers or professionals could generate a bias in the predictions. Background and related work are given in A.1 and A.2.

2 DATA AND EXPERIMENTS

The data is the record of ambulance departures from the FD of Doubs, in France, from 2018 to 2021. Each sample represents a departure with the attributes described in A.3-Table 1-Fig. 3. Originally, the target variable is the ambulance's engagement time in minutes, i.e., from alert to return. This is converted into a binary indicator (*iTime*) for classification, where class 1 represents an engagement time greater than the average of the last 3 hours, and class 0 otherwise. Here, 4 sensitive variables are identified: *type* (sanitized with GRR Kairouz et al. (2016); Arcolezi et al. (2022a) to protect victim's privacy); *nbP* and *nbV* (sanitized with Geometric Ghosh et al. (2008); Bellamy et al. (2018) to protect fireman's privacy); and *iProf*, where 0 (unprivileged group) means *nbV*>*nbP* and 1 (privileged group) otherwise. These groups could naturally generate a bias in the model, since there are more volunteers with basic training and few with full training that can result in more and less engagement time, respectively. Besides, with DP, the bias effect could increase Fioretto et al. (2022); Pujol et al. (2020), hence, we will use Rew to correct it. For *type*, *nbP*, and *nbV*, ϵ will be divided uniformly.

For modeling, we will use RandomForestClassifier. The training and testing sets are departures from 01/01/2018 to 30/06/2021 and from 01/07/2021 to 31/12/2021, respectively. Our experiments will be: **E1)** Recover the best base model and configuration with Bayesian optimization Bergstra et al. (2013). The best configuration will be reused in the next models. **E2)** Correct the bias according to *iProf* and with Rew. **E3)** Assess privacy's level ($\forall \epsilon \in L = \{0.1, 0.2, ...1, 2, ..., 15\}$) and fairness of the classifier. Since DP algorithms are randomized, we will report average results over 20 runs.

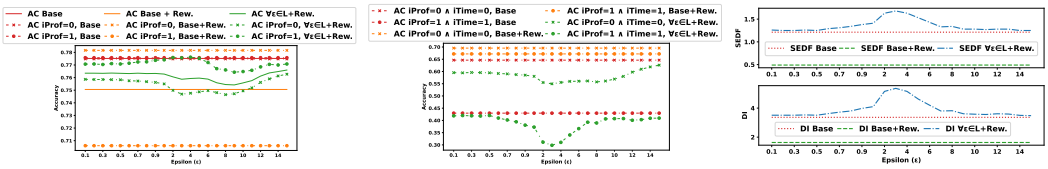

(a) (b) (c)

Figure 1: Models performance.

3 Results and Discussion

We used Accuracy (AC) and F1-Score (F1) to assess model performance, and Disparate Impact (DI) and Smoothed Empirical Differential Fairness (SEDF) to measure their fairness. In A.3-Table 2, modeling settings for E1 are described. In A.3-Table 3, results for E1, E2, and E3 are presented.

In Fig. 1a, 2 levels of AC are analyzed. The first level is the general AC of the models (line), where private models present a better result than the fair base model but lower than the base. The second level is the AC by category of *iProf* (dashdotted), where results with $1 < \epsilon \leq 15$ degrade and $0.1 \leq \epsilon \leq 1$ stabilize. Why does this stabilization occur? Fig. 2 answers this question by analyzing feature importance for 4 values of ϵ+Rew. Here, we see that the sensitive variables *nbP* and *nbV* reduce their impact as ϵ decreases, i.e., the stabilization occurs because the classifier uses them less to build trees and increases the importance of others such as *center*.

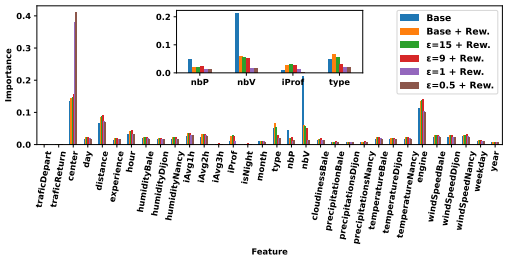

Figure 2: Feature importance.

In Fig. 1b, the third level of AC is analyzed, i.e., by category of iProf and iTime. Here, Base+Rew shows the best performance and the AC of the unprivileged group in DP models decreases rapidly from $\epsilon = 15$ until $\epsilon = 1$ unlike the privileged group and compared to the base model groups. This is because *nbV* lost its importance as shown in Fig. 2. In Fig. 1c, fairness metrics SEDF and DI are analyzed, where it is verified that the models based on sensitive tabular variables preprocessed with Geometric, GRR, and Rew mechanisms are less fair than the Base and Base+Rew models. This has also been observed in Bagdasaryan et al. (2019); Fioretto et al. (2022); Ganev et al. (2022), where DP was applied during the inprocessing stage. However, as discussed in de Oliveira et al. (2023), in real life, an optimal hyperparameter search should be performed since data has been modified with DP. This would consume more resources, but a better model would be recovered for production deployment.

4 Conclusion

This study developed a model to identify if the ambulance will spend more time engaged, and consequently, to reassign ambulances to the affected center in a short time. We determined that it is possible to keep agents' and victims' privacy, recovering an ML model with good accuracy but less fair since the negative bias towards the underprivileged group was greater than the baselines. This is explained because we kept the same hyperparameter configuration for all models and varied data according to their preprocessing with GRR, Geometric, and Rew. The results showed: a) Good AC in the first level, but a degradation of the AC in the unprivileged group. b) For a value of $\epsilon \leq 1$, a stabilization of performance was observed since the sensitive variable (*nbV*) lost its importance and the classifier replaced it with another one, such as *center*. Hence, some future directions are: to analyze the feature importance in private models and assess the DP impact using different encoding for sensitive variables (label, one hot encoding, etc). Also, another study would be to determine if the same results (performance and feature impact) would be obtained using different types of classifiers (linear, boosting, neural network, etc).

URM STATEMENT

The authors acknowledge that at least one key author of this work meets the URM criteria of ICLR 2023 Tiny Papers Track.

REFERENCES

Ankit Agrawal, Sophia J. Abraham, Benjamin Burger, Chichi Christine, Luke Fraser, John M. Hoeksema, Sarah Hwang, Elizabeth Travnik, Shreya Kumar, Walter Scheirer, Jane Cleland-Huang, Michael Vierhauser, Ryan Bauer, and Steve Cox. The next generation of human-drone partnerships: Co-designing an emergency response system. In *Proceedings of the 2020 CHI Conference on Human Factors in Computing Systems*. ACM, April 2020. doi: 10.1145/3313831.3376825.

Héber H. Arcolezi, Jean-François Couchot, Sébastien Gambs, Catuscia Palamidessi, and Majid Zolfaghari. Multi-freq-LDPy: Multiple frequency estimation under local differential privacy in python. In *Computer Security – ESORICS 2022*, pp. 770–775. Springer Nature Switzerland, 2022a. doi: 10.1007/978-3-031-17143-7_40.

Héber H. Arcolezi, Selene Cerna, Jean-François Couchot, Christophe Guyeux, and Abdallah Makhoul. Privacy-preserving prediction of victim's mortality and their need for transportation to health facilities. *IEEE Transactions on Industrial Informatics*, 18(8):5592–5599, 2022b. doi: 10.1109/TII.2021.3123588.

Eugene Bagdasaryan, Omid Poursaeed, and Vitaly Shmatikov. *Differential Privacy Has Disparate Impact on Model Accuracy*. Curran Associates Inc., Red Hook, NY, USA, 2019.

Rachel K. E. Bellamy, Kuntal Dey, Michael Hind, Samuel C. Hoffman, Stephanie Houde, Kalapriya Kannan, Pranay Lohia, Jacquelyn Martino, Sameep Mehta, Aleksandra Mojsilovic, Seema Nagar, Karthikeyan Natesan Ramamurthy, John Richards, Diptikalyan Saha, Prasanna Sattigeri, Moninder Singh, Kush R. Varshney, and Yunfeng Zhang. AI Fairness 360: An extensible toolkit for detecting, understanding, and mitigating unwanted algorithmic bias, October 2018.

J. Bergstra, D. Yamins, and D. D. Cox. Making a science of model search: Hyperparameter optimization in hundreds of dimensions for vision architectures. In *Proceedings of the 30th International Conference on International Conference on Machine Learning*, ICML'13, pp. I–115–I–123. JMLR, 2013.

Bison-Futé. Les prévisions de trafic. Available online: `https://www.bison-fute.gouv.fr` (accessed on February 27, 2023).

Leo Breiman. Random forests. *Machine Learning*, 45(1):5–32, 2001. doi: 10.1023/a:1010933404324.

Toon Calders, Faisal Kamiran, and Mykola Pechenizkiy. Building classifiers with independency constraints. In *2009 IEEE International Conference on Data Mining Workshops*, pp. 13–18, 2009. doi: 10.1109/ICDMW.2009.83.

Selene Cerna, Héber H. Arcolezi, Christophe Guyeux, Guillaume Royer-Fey, and Céline Chevallier. Machine learning-based forecasting of firemen ambulances' turnaround time in hospitals, considering the COVID-19 impact. *Applied Soft Computing*, 109:107561, September 2021. doi: 10.1016/j.asoc.2021.107561.

Anderson Santana de Oliveira, Caelin Kaplan, Khawla Mallat, and Tanmay Chakraborty. An empirical analysis of fairness notions under differential privacy. PPAI 2023, 4th AAAI Workshop on Privacy-Preserving Artificial Intelligence, 13 February 2023, Washington DC, USA / Also Submitted to ArXiV, 6 February 2023, 2023. © AAAI. Personal use of this material is permitted. The definitive version of this paper was published in PPAI 2023, 4th AAAI Workshop on Privacy-Preserving Artificial Intelligence, 13 February 2023, Washington DC, USA / Also Submitted to ArXiV, 6 February 2023 and is available at :.

Cynthia Dwork. Differential privacy. In *Automata, Languages and Programming*, pp. 1–12. Springer Berlin Heidelberg, 2006. doi: 10.1007/11787006_1.

Ferdinando Fioretto, Cuong Tran, Pascal Van Hentenryck, and Keyu Zhu. Differential Privacy and Fairness in Decisions and Learning Tasks: A Survey. *arXiv e-prints*, art. arXiv:2202.08187, February 2022.

James R. Foulds, Rashidul Islam, Kamrun Naher Keya, and Shimei Pan. An intersectional definition of fairness. In *2020 IEEE 36th International Conference on Data Engineering (ICDE)*, pp. 1918–1921, 2020. doi: 10.1109/ICDE48307.2020.00203.

Georgi Ganev, Bristena Oprisanu, and Emiliano De Cristofaro. Robin hood and matthew effects: Differential privacy has disparate impact on synthetic data. In *International Conference on Machine Learning*, pp. 6944–6959. PMLR, 2022.

Arpita Ghosh, Tim Roughgarden, and Mukund Sundararajan. Universally utility-maximizing privacy mechanisms, 2008.

Christophe Guyeux. Predictions in pre-hospital emergency transport in france: A state of the art. In *Signal Processing and Vision*. Academy and Industry Research Collaboration Center (AIRCC), December 2022. doi: 10.5121/csit.2022.122205.

Peter Kairouz, Keith Bonawitz, and Daniel Ramage. Discrete distribution estimation under local privacy. In *International Conference on Machine Learning*, pp. 2436–2444. PMLR, 2016.

Faisal Kamiran and Toon Calders. Data preprocessing techniques for classification without discrimination. *Knowledge and Information Systems*, 33:1 – 33, 2011.

Xu Lian, Sarah Melancon, Jon-Ross Presta, Adam Reevesman, Brian Spiering, and Diane Woodbridge. Scalable real-time prediction and analysis of san francisco fire department response times. In *2019 IEEE SmartWorld, Ubiquitous Intelligence & Computing, Advanced & Trusted Computing, Scalable Computing & Communications, Cloud & Big Data Computing, Internet of People and Smart City Innovation (SmartWorld/SCALCOM/UIC/ATC/CBDCom/IOP/SCI)*. IEEE, August 2019. doi: 10.1109/smartworld-uic-atc-scalcom-iop-sci.2019.00154.

Météo-France. Service météorologique national français. Available online: `https://donneespubliques.meteofrance.fr/?fond=produit&id_produit=90&id_rubrique=32` (accessed on February 27, 2023).

Thiago Fonseca Morello, Rossano Marchetti Ramos, Liana O. Anderson, Nathan Owen, Thais Michele Rosan, and Lara Steil. Predicting fires for policy making: Improving accuracy of fire brigade allocation in the brazilian amazon. *Ecological Economics*, 169:106501, March 2020. doi: 10.1016/j.ecolecon.2019.106501.

Aileen Nielsen. *Practical fairness*. O'Reilly Media, Sebastopol, CA, December 2020.

F. Pedregosa, G. Varoquaux, A. Gramfort, V. Michel, B. Thirion, O. Grisel, M. Blondel, P. Prettenhofer, R. Weiss, V. Dubourg, J. Vanderplas, A. Passos, D. Cournapeau, M. Brucher, M. Perrot, and E. Duchesnay. Scikit-learn: Machine learning in Python. *Journal of Machine Learning Research*, 12:2825–2830, 2011.

David Pujol, Ryan McKenna, Satya Kuppam, Michael Hay, Ashwin Machanavajjhala, and Gerome Miklau. Fair decision making using privacy-protected data. In *Proceedings of the 2020 Conference on Fairness, Accountability, and Transparency*, FAT* '20, pp. 189–199, New York, NY, USA, 2020. Association for Computing Machinery. ISBN 9781450369367. doi: 10.1145/3351095.3372872.

Julian Weidinger. What is known and what remains unexplored: A review of the firefighter information technologies literature. *International Journal of Disaster Risk Reduction*, 78:103115, 2022. ISSN 2212-4209. doi: https://doi.org/10.1016/j.ijdrr.2022.103115.

Tianqing Zhu, Gang Li, Wanlei Zhou, and Philip S Yu. *Differential privacy and applications*. Advances in Information Security. Springer International Publishing, Basel, Switzerland, 1 edition, August 2017.

A APPENDIX

A.1 BACKGROUND

In this section, basic notions related to ML, DP and Fairness techniques, that we use, are given.

Supervised Learning. A dataset D contains pairs (x_i, y_i) of input and true label, respectively, and a function g outputs a predicted label $g(x_i)$. The training process of g is done to capture patterns from x_i while minimizing the loss function $\mathcal{L}(g(x_i), y_i)$ until obtaining the best model g. An example of technique is **Random Forest**, an ensemble learning method of type bagging Breiman (2001); Pedregosa et al. (2011).

Differential Privacy. A randomized mechanism M provides (ϵ, δ)-DP for every set of outputs S, and for any neighbouring datasets D and $D^{'}$ differing in one element, if and only if M satisfies $Pr[M(D) \in S] \leq e^{\epsilon} Pr[M(D^{'}) \in S] + \delta$. The privacy budget is represented by $\epsilon \geq 0$, the smaller the more privacy. If $\delta = 0$, M gives ϵ-DP in strict definition, and $\delta > 0$ gives freedom to violate ϵ-DP for some low probability events. For multi-step mechanisms $M = \{M_1, M_2, ..., M_m\}$, 3 properties must be considered: *a) Post-processing*, where any additional processing to M's output or the composition $M_1(M_2(.))$ still satisfies ϵ-DP; *b) Sequential composition*, where mechanisms M, sequentially performed on a dataset with their respective ϵ_i-DP, will provide $(\sum_{i=1}^{m} \epsilon_i)$-DP; and, *c) Parallel composition*, where mechanisms M, applied on a disjointed subset of the entire dataset and with their respective ϵ_i-DP, will provide $(max\{\epsilon_1, ..., \epsilon_m\})$-DP Dwork (2006); Zhu et al. (2017).

Geometric. It is a DP mechanism, defined as $M(D^{'}) = f(D^{'}) + z$, where $f(D^{'})$ is the true query result and z is a random variable from the double geometric distribution $\frac{1-e^{-\epsilon\Delta^{-1}}}{1+e^{-\epsilon\Delta^{-1}}} e^{-\epsilon|z|\Delta^{-1}}$. For counting functions, the sensitivity (Δ) is 1, which will be our case Ghosh et al. (2008); Bellamy et al. (2018).

Direct Encoding. Also known as Generalized Randomized Response (GRR), that extends the Randomized Response DP mechanism for cases where $k \geq 2$, being k the size of a set of categories $A = \{v_1, ..., v_k\}$ of a given attribute. Thus, GRR(v) outputs the true value with probability p, and any other value $v^{'} \in A$ such that $v^{'} \neq v$ with probability $1 - p$. More precisely, the perturbation is defined as a conditional function: if $y = v$, then $p = \frac{e^{\epsilon}}{e^{\epsilon}+k-1}$; or if $y \neq v$, then $q = \frac{1}{e^{\epsilon}+k-1}$, $\forall y \in A$ Kairouz et al. (2016); Arcolezi et al. (2022a).

Reweighing. It is a mechanism that considers independency constraint to build a fair classifier. Let's say a sample x has a binary protected attribute B with value b ($b \in dom(B) = \{0, 1\}$). If its expected probability (P_{exp}) is higher than its actual probability (P_{act}), the bias will be towards the negative class (class $-$) for $B = b$ and the weights will be assigned to b with respect to positive class (class $+$): $W(B = b|x(Class) = +) = \frac{P_{exp}(b \wedge +)}{P_{act}(b \wedge +)}$. The weight of b for class $-$ will be: $W(B = b|x(Class) = -) = \frac{P_{exp}(b \wedge -)}{P_{act}(b \wedge -)}$. The same calculation is applied for weights of $\bar{b}$ for classes $+$ and $-$ Calders et al. (2009); Kamiran & Calders (2011); Bellamy et al. (2018).

Disparate Impact. It is a fairness metric, which measures the positive proportions in the predicted outcomes. Values less than and greater than 1 indicate positive and negative bias, respectively. The ideal value should be close to 1 Nielsen (2020).

Smoothed Empirical Differential Fairness. It is a fairness metric, which compares the differential of smoothed probability between groups and provides a closely related privacy guarantee to differential privacy. The ideal value should be close to 0 Foulds et al. (2020); Bellamy et al. (2018).

A.2 RELATED WORK

In recent years, artificial intelligence (AI) has allowed the accelerated development of various sectors of society such as health, education, environment, etc. However, in areas such as civil security, there is still much to be developed. Here, for example, we find emergency medical services and fire departments, which have been recording their interventions for years. These records are a very rich source to exploit with AI. The main problem for these organizations is that as the population grows, the number of incidents also grows (cardiac arrests, suicides, drownings, traffic accidents, etc.) and

if the resources like ambulances to rescue people are not enough, a breakdown in the system would occur and the consequences could be human, economic, and material losses. For this reason, the present study develops an intelligent model for predicting ambulance engagement time, defined as an indicator. This indicator allows to identify if the engagement time of an ambulance will be higher than the average of all ambulance times in the last 3 hours. In this way, time flows can be perceived during the day, dynamically reorganize ambulances between nearby centers, avoid breakdowns in the system, and keep people protected. When reviewing the literature and to the author's knowledge, we found articles on the optimization, simulation, and prediction of some firefighters operations Guyeux (2022); Arcolezi et al. (2022b); Cerna et al. (2021); Morello et al. (2020); Agrawal et al. (2020); Lian et al. (2019) and a gap analysis on firefighter information technologies literature Weidinger (2022). However, none of them developed a predictive model for ambulance engagement time. In addition, we complete our study by applying DP mechanisms to preserve the privacy of victims and agents. Since the application of DP mechanisms may generate or increase an existing bias Fioretto et al. (2022); Pujol et al. (2020), we also analyze two fairness metrics in our private models, which have not yet been explored in this area.

A.3 Complementary configurations and results

Table 1: Description of the variables considered.

Type	Variable	Description	Values
Input - Internal	center	Fire station from which the ambulance departed.	{1,2,...,71}
	distance	Distance between the fire station and the incident scene (km).	[0.001-7103]
	engine	Type of vehicle (ambulance) used.	{1,2,3,4,5}
	experience	Number of agent departures in the last month.	{0,1,...,108}
	nbP	Number of professional firefighters.	{1,2,3,4}
	nbV	Number of volunteer firefighters.	{1,2,3,4,5}
	iProf	Crew professionalism indicator, where: iProf=1 means nbV≤nbP, iProf=0 means nbV>nbP.	{0,1}
	type	Type of intervention (altered consciousness, cardio-respiratory arrest, imminent delivery, etc).	{0,1,...,36}
	day	Day of the intervention.	{1,2,...,31}
	hour	Hour of the intervention.	{0,1,...,23}
	iAvg1h	Indicator for the average engagement time, where: iAvg1h=1 means that the average engagement time of all ambulance departures in the last hour is greater than the average of all those made in the last month, iAvg1h=0 means otherwise.	{0,1}
	iAvg2h	Indicator for the average engagement time, where: iAvg1h=1 means that the average engagement time of all ambulance departures in the last 2 hours is greater than the average of all those made in the last month, iAvg1h=0 means otherwise.	{0,1}
	iAvg3h	Indicator for the average engagement time, where: iAvg1h=1 means that the average engagement time of all ambulance departures in the last 3 hours is greater than the average of all those made in the last month, iAvg1h=0 means otherwise.	{0,1}
	isNight	Indicator to identify if it is night or day.	{0,1}
	month	Month of the intervention.	{1,2,...,12}
	weekday	Weekday of the intervention.	{0,1,...,6}
	year	Year of the intervention.	{2018, 2019, 2020, 2021}
Input - Traffic (Bison-Futé)	traficDepart	Traffic level indicator on exit roads.	{0,1,2,3}
	traficReturn	Traffic level indicator on arrival roads.	{0,1,2,3}
Input - Weather (Météo-France)	cloudinessBale	Cloudiness registered by an station in Bâle.	[0-101]
	humidityBale	Humidity registered by an station in Bâle.	[18-100]
	humidityDijon	Humidity registered by an station in Dijon.	[1-100]
	humidityNancy	Humidity registered by an station in Nancy.	[14-100]
	precipitationBale	Precipitation registered by an station in Bâle.	[-0.1-22.1]
	precipitationsDijon	Precipitation registered by an station in Dijon.	[-0.1-24.3]
	precipitationsNancy	Precipitation registered by an station in Nancy.	[-0.1-8.8]
	temperatureBale	Temperature registered by an station in Bâle.	[257.05-311.45]
	temperatureDijon	Temperature registered by an station in Dijon.	[262.55-312.05]
	temperatureNancy	Temperature registered by an station in Nancy.	[261.15-312.55]
	windSpeedBale	Wind speed registered by an station in Bâle.	[0.0-15.1]
	windSpeedDijon	Wind speed registered by an station in Dijon.	[0.0-16.6]
	windSpeedNancy	Wind speed registered by an station in Nancy.	[0.0-16.9]
Output - Target	iTime	Indicator for ambulance engagement time, where: $iTime = 1$ means that the ambulance engagement time is greater than the average engagement time of all ambulances in the last 3 hours, $iTime = 0$ means otherwise	{0,1}

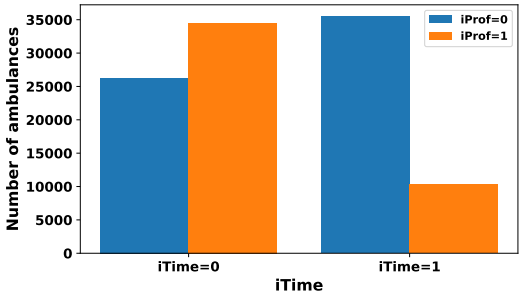

Figure 3: Distribution of the target variable by privileged and unprivileged group.

Table 2: Definition of the search space for RandomForestClassifier, the best hyperparameter configuration recovered, and the loss function and number of iterations for Bayesian Optimization.

Search space	**Best configuration**
n_estimators: [200-600]	200
max_depth: [10-15]	14
max_features: [0.5-1]	0.8
max_samples: [0.5-1]	0.95
class_weight: {0:1-5, 1:1-5}	{0:1, 1:2}
Loss function	$\mathcal{L}$ = (AC iTime=0)x(AC iTime=1), if AC iTime=0 $>$ 0.6 $\wedge$ AC iTime=1 $>$ 0.6 $\mathcal{L}$ = (AC iTime=0)+(AC iTime=1), otherwise.
Number of iterations	50

Table 3: Results of fairness and performance measures for experiments E1, E2, E3.

Model	**Metrics**					
	AC	**AC iProf=0**	**AC iProf=1**	**F1**	**DI**	**SEDF**
E1 - Base	0.78	0.77 $\wedge$ iTime=0: 0.65 $\wedge$ iTime=1: 0.88	0.78 $\wedge$ iTime=0: 0.88 $\wedge$ iTime=1: 0.43	0.77	3.36	1.21
E2 - Base + Rew.	0.75	0.78 $\wedge$ iTime=0: 0.70 $\wedge$ iTime=1: 0.85	0.71 $\wedge$ iTime=0: 0.72 $\wedge$ iTime=1: 0.67	0.75	1.63	0.49
E3 - $\epsilon = 15$ + Rew	0.77	0.76 $\wedge$ iTime=0: 0.62 $\wedge$ iTime=1: 0.87	0.77 $\wedge$ iTime=0: 0.89 $\wedge$ iTime=1: 0.39	0.76	3.71	1.31
E3 - $\epsilon = 9$ + Rew	0.75	0.74 $\wedge$ iTime=0: 0.55 $\wedge$ iTime=1: 0.90	0.77 $\wedge$ iTime=0: 0.88 $\wedge$ iTime=1: 0.41	0.75	3.70	1.30
E3 - $\epsilon = 1$ + Rew	0.76	0.76 $\wedge$ iTime=0: 0.58 $\wedge$ iTime=1: 0.89	0.77 $\wedge$ iTime=0: 0.89 $\wedge$ iTime=1: 0.37	0.76	4.07	1.40
E3 - $\epsilon = 0.5$ + Rew	0.76	0.76 $\wedge$ iTime=0: 0.59 $\wedge$ iTime=1: 0.89	0.77 $\wedge$ iTime=0: 0.88 $\wedge$ iTime=1: 0.42	0.76	3.58	1.28