Reliable and fair machine learning for risk assessment

Pereira Barata, A.P.

References

[Abellán and Masegosa, 2010] Abellán, J. and Masegosa, A. R. (2010). Bagging decision trees on data sets with classification noise. In *International Symposium on Foundations of Information and Knowledge Systems*, pages 248–265. Springer.

[Abellán and Moral, 2003] Abellán, J. and Moral, S. (2003). Building classification trees using the total uncertainty criterion. *International Journal of Intelligent Systems*, 18(12):1215–1225.

[Ahirwar, 2020] Ahirwar, J. P. (2020). Five laws of library science and information economics. *Informatics Studies*, 7(1).

[Alazzam et al., 2019] Alazzam, H., Alsmady, A., and Shorman, A. A. (2019). Supervised detection of IoT botnet attacks. In *Proceedings of the Second International Conference on Data Science, E-Learning and Information Systems*, pages 1–6.

[Alcalá-Fdez et al., 2009] Alcalá-Fdez, J., Sanchez, L., Garcia, S., del Jesus, M. J., Ventura, S., Garrell, J. M., Otero, J., Romero, C., Bacardit, J., Rivas, V. M., et al. (2009). KEEL: a software tool to assess evolutionary algorithms for data mining problems. *Soft Computing*, 13(3):307–318.

[Amer and Goldstein, 2012] Amer, M. and Goldstein, M. (2012). Nearest-neighbor and clustering based anomaly detection algorithms for rapidminer. In *Proceedings of the 3rd RapidMiner Community Meeting and Conference (RCOMM 2012)*, pages 1–12.

[Amiri and Jensen, 2016] Amiri, M. and Jensen, R. (2016). Missing data imputation using fuzzy-rough methods. *Neurocomputing*, 205:152–164.

[Angenent et al., 2020] Angenent, M. N., Pereira Barata, A., and Takes, F. W. (2020). Large-scale machine learning for business sector prediction. In *Proceedings of the 35th Annual ACM Symposium on Applied Computing*, pages 1143–1146.

[Angluin and Laird, 1988] Angluin, D. and Laird, P. (1988). Learning from noisy examples. *Machine Learning*, 2(4):343–370.

[Asuncion and Newman, 2007] Asuncion, A. and Newman, D. (2007). UCI machine learning repository.

[Azar and El-Metwally, 2013] Azar, A. T. and El-Metwally, S. M. (2013). Decision tree classifiers for automated medical diagnosis. *Neural Computing and Applications*, 23(7):2387–2403.

[Barabási, 2016] Barabási, A. L. (2016). *Network Science*. Cambridge University Press.

[Barocas et al., 2017] Barocas, S., Hardt, M., and Narayanan, A. (2017). Fairness in machine learning. *Nips tutorial*, 1:2.

[Bartlett et al., 2006] Bartlett, P. L., Jordan, M. I., and McAuliffe, J. D. (2006). Convexity, classification, and risk bounds. *Journal of the American Statistical Association*, 101(473):138–156.

[Beale and Little, 1975] Beale, E. M. and Little, R. J. (1975). Missing values in multivariate analysis. *Journal of the Royal Statistical Society: Series B (Methodological)*, 37(1):129–145.

[Bechavod and Ligett, 2017] Bechavod, Y. and Ligett, K. (2017). Penalizing unfairness in binary classification. *arXiv preprint arXiv:1707.00044*.

[Beigman Klebanov and Beigman, 2009] Beigman Klebanov, B. and Beigman, E. (2009). From annotator agreement to noise models. *Computational Linguistics*, 35(4):495–503.

[Bennett, 2001] Bennett, D. A. (2001). How can I deal with missing data in my study? *Australian and New Zealand Journal of Public Health*, 25(5):464–469.

[Bennett, 2000] Bennett, P. N. (2000). Assessing the calibration of naive bayes posterior estimates. Technical report, Carnegie-Mellon Univ Pittsburgh PA School of Computer Science.

[Bergstra and Bengio, 2012] Bergstra, J. and Bengio, Y. (2012). Random search for hyper-parameter optimization. *Journal of Machine Learning Research*, 13(2).

[Bertsimas et al., 2017] Bertsimas, D., Pawlowski, C., and Zhuo, Y. D. (2017). From predictive methods to missing data imputation: an optimization approach. *The Journal of Machine Learning Research*, 18(1):7133–7171.

[Böken, 2021] Böken, B. (2021). On the appropriateness of Platt scaling in classifier calibration. *Information Systems*, 95:101641.

[Bonner et al., 2015] Bonner, S., McGough, A. S., Kureshi, I., Brennan, J., Theodoropoulos, G., Moss, L., Corsar, D., and Antoniou, G. (2015). Data quality assessment and anomaly detection via map/reduce and linked data: a case study in the medical domain. In *2015 IEEE International Conference on Big Data (Big Data)*, pages 737–746. IEEE.

[Breiman, 2001] Breiman, L. (2001). Random forests. *Machine Learning*, 45(1):5–32.

[Breiman et al., 1984] Breiman, L., Friedman, J., Olshen, R., and Stone, C. (1984). *Classification and Regression Trees*. Wadsworth.

[Breunig et al., 2000] Breunig, M. M., Kriegel, H.-P., Ng, R. T., and Sander, J. (2000). LOF: identifying density-based local outliers. In *Proceedings of the 2000 ACM SIGMOD International Conference on Management of Data*, pages 93–104.

[Brink et al., 2016] Brink, H., Richards, J., and Fetherolf, M. (2016). *Real-world machine learning*. Simon and Schuster.

[Brodley and Friedl, 1996] Brodley, C. E. and Friedl, M. A. (1996). Improving automated land cover mapping by identifying and eliminating mislabeled observations from training data. In *IGARSS'96, 1996 International Geoscience and Remote Sensing Symposium*, volume 2, pages 1379–1381. IEEE.

[Brodley and Friedl, 1999] Brodley, C. E. and Friedl, M. A. (1999). Identifying mislabeled training data. *Journal of Artificial Intelligence Research*, 11:131–167.

[Buolamwini and Gebru, 2018] Buolamwini, J. and Gebru, T. (2018). Gender shades: intersectional accuracy disparities in commercial gender classification. In *Conference on Fairness, Accountability and Transparency*, pages 77–91. PMLR.

[Cariou and Wolff, 2011] Cariou, P. and Wolff, F.-C. (2011). Do port state control inspections influence flag-and class-hopping phenomena in shipping? *Journal of Transport Economics and Policy (JTEP)*, 45(2):155–177.

[Carpenter and Kenward, 2012] Carpenter, J. and Kenward, M. (2012). *Multiple Imputation and its Application*. John Wiley & Sons.

[Chan and Treleaven, 2015] Chan, S. and Treleaven, P. (2015). Continuous model selection for large-scale recommender systems. In *Handbook of Statistics*, volume 33, pages 107–124. Elsevier.

[Chapelle et al., 2006] Chapelle, O., Scholkopf, B., and Zien, A. (2006). Semi-supervised learning. 2006. *Cambridge, Massachusettes: The MIT Press View Article*, 2.

[Chen et al., 2014] Chen, M., Yang, C., Li, C., Hou, L., Chen, X., and Zhao, H. (2014). Admixture mapping analysis in the context of GWAS with GAW18 data. In *BMC Proceedings*, pages 1–5. BioMed Central.

[Chen and Guestrin, 2016] Chen, T. and Guestrin, C. (2016). XGBoost: a scalable tree boosting system. In *Proceedings of the 22nd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, pages 785–794.

[Cho et al., 2020] Cho, J., Hwang, G., and Suh, C. (2020). A fair classifier using kernel density estimation. *Advances in Neural Information Processing Systems*, 33:15088–15099.

[Choi et al., 2019] Choi, J., Dekkers, O. M., and le Cessie, S. (2019). A comparison of different methods to handle missing data in the context of propensity score analysis. *European Journal of Epidemiology*, 34(1):23–36.

[Choudhary and Gianey, 2017] Choudhary, R. and Gianey, H. K. (2017). Comprehensive review on supervised machine learning algorithms. In *2017 International Conference on Machine Learning and Data Science (MLDS)*, pages 37–43. IEEE.

[Chu et al., 2013] Chu, X., Ilyas, I. F., and Papotti, P. (2013). Discovering denial constraints. *Proceedings of the VLDB Endowment*, 6(13):1498–1509.

[Claesen and De Moor, 2015] Claesen, M. and De Moor, B. (2015). Hyperparameter search in machine learning. *arXiv preprint arXiv:1502.02127*.

[CLAIRE, 2021] CLAIRE (2021). Response to the European Commission's proposal for AI regulation and 2021 coordinated plan on AI. `https://claire-ai.org/wp-content/uploads/2021/08/CLAIRE-EC-AI-Regulation-Feedback.pdf`.

[Corbett-Davies and Goel, 2018] Corbett-Davies, S. and Goel, S. (2018). The measure and mismeasure of fairness: a critical review of fair machine learning. *arXiv preprint arXiv:1808.00023*.

[Corbett-Davies et al., 2017] Corbett-Davies, S., Pierson, E., Feller, A., Goel, S., and Huq, A. (2017). Algorithmic decision making and the cost of fairness. In *Proceedings of the 23rd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, pages 797–806.

[Cuendet et al., 2007] Cuendet, S., Hakkani-Tür, D., and Shriberg, E. (2007). Automatic labeling inconsistencies detection and correction for sentence unit segmentation in conversational speech. In *International Workshop on Machine Learning for Multimodal Interaction*, pages 144–155. Springer.

[Dardanoni et al., 2011] Dardanoni, V., Modica, S., and Peracchi, F. (2011). Regression with imputed covariates: a generalized missing-indicator approach. *Journal of Econometrics*, 162(2):362–368.

[Dastian, 2018] Dastian, J. (2018). Amazon scraps secret ai recruiting tool that showed bias against women. *Reuters*.

[de Bruin et al., 2022] de Bruin, G. J., Pereira Barata, A., van den Herik, H. J., Takes, F. W., and Veenman, C. J. (2022). Fair automated assessment of non-compliance in cargo ship networks. *EPJ Data Science*, 11(1):13.

[Diekmann and Jann, 2010] Diekmann, A. and Jann, B. (2010). Benford's law and fraud detection: facts and legends. *German Economic Review*, 11(3):397–401.

[Dietterich, 2000] Dietterich, T. G. (2000). An experimental comparison of three methods for constructing ensembles of decision trees: bagging, boosting, and randomization. *Machine Learning*, 40(2):139–157.

[Ding and Simonoff, 2010] Ding, Y. and Simonoff, J. S. (2010). An investigation of missing data methods for classification trees applied to binary response data. *Journal of Machine Learning Research*, 11(1).

[Dogru and Subasi, 2018] Dogru, N. and Subasi, A. (2018). Traffic accident detection using random forest classifier. In *2018 15th Learning and Technology Conference (L&T)*, pages 40–45. IEEE.

[Dressel and Farid, 2018] Dressel, J. and Farid, H. (2018). The accuracy, fairness, and limits of predicting recidivism. *Science Advances*, 4(1):eaao5580.

[Dwork et al., 2012] Dwork, C., Hardt, M., Pitassi, T., Reingold, O., and Zemel, R. (2012). Fairness through awareness. In *Proceedings of the 3rd Innovations in Theoretical Computer Science Conference*, pages 214–226.

[Enders, 2010] Enders, C. K. (2010). *Applied Missing Data Analysis*. Guilford Press.

[European Commission, 2018a] European Commission (2018a). Waste classification. `https://ec.europa.eu/environment/waste/framework/list.htm`.

[European Commission, 2018b] European Commission (2018b). Waste legislation. https://ec.europa.eu/environment/waste/legislation.

[European Commission, 2019a] European Commission (2019a). A definition of AI: main capabilities and disciplines. https://ec.europa.eu/newsroom/dae/document.cfm?doc_id=56341.

[European Commission, 2019b] European Commission (2019b). Ethics guidelines for trustworthy AI. https://ec.europa.eu/newsroom/dae/document.cfm?doc_id=60419.

[European Commission, 2019c] European Commission (2019c). Proposal for a regulation on a European approach for artificial intelligence. https://digital-strategy.ec.europa.eu/en/library/proposal-regulation-european-approach-artificial-intelligence.

[European Commission, 2021] European Commission (2021). Regulatory framework proposal on artificial intelligence. https://digital-strategy.ec.europa.eu/en/policies/regulatory-framework-ai.

[Feldman et al., 2015] Feldman, M., Friedler, S. A., Moeller, J., Scheidegger, C., and Venkatasubramanian, S. (2015). Certifying and removing disparate impact. In *Proceedings of the 21th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, pages 259–268.

[Feng et al., 2011] Feng, X., Wu, S., and Liu, Y. (2011). Imputing missing values for mixed numeric and categorical attributes based on incomplete data hierarchical clustering. In *International Conference on Knowledge Science, Engineering and Management*, pages 414–424. Springer.

[Flach, 2016] Flach, P. A. (2016). ROC analysis. In *Encyclopedia of Machine Learning and Data Mining*, pages 1–8. Springer.

[Frénay and Verleysen, 2013] Frénay, B. and Verleysen, M. (2013). Classification in the presence of label noise: a survey. *IEEE Transactions on Neural Networks and Learning Systems*, 25(5):845–869.

[Frery et al., 2017] Frery, J., Habrard, A., Sebban, M., Caelen, O., and He-Guelton, L. (2017). Efficient top rank optimization with gradient boosting for supervised anomaly detection. In *Joint European Conference on Machine Learning and Knowledge Discovery in Databases*, pages 20–35. Springer.

[Friedman, 2001] Friedman, J. H. (2001). Greedy function approximation: a gradient boosting machine. *Annals of Statistics*, pages 1189–1232.

[Fürber, 2016] Fürber, C. (2016). Semantic technologies. In *Data Quality Management with Semantic Technologies*, pages 56–68. Springer.

[Gamberger et al., 1996] Gamberger, D., Lavrač, N., and Džeroski, S. (1996). Noise elimination in inductive concept learning: a case study in medical diagnosis. In *International Workshop on Algorithmic Learning Theory*, pages 199–212. Springer.

[Gamberger et al., 1999] Gamberger, D., Lavrac, N., and Groselj, C. (1999). Experiments with noise filtering in a medical domain. In *ICML*, volume 99, pages 143–151.

[García-Laencina et al., 2015] García-Laencina, P. J., Abreu, P. H., Abreu, M. H., and Afonoso, N. (2015). Missing data imputation on the 5-year survival prediction of breast cancer patients with unknown discrete values. *Computers in Biology and Medicine*, 59:125–133.

[García-Laencina et al., 2010] García-Laencina, P. J., Sancho-Gómez, J.-L., and Figueiras-Vidal, A. R. (2010). Pattern classification with missing data: a review. *Neural Computing and Applications*, 19(2):263–282.

[Garciarena and Santana, 2017] Garciarena, U. and Santana, R. (2017). An extensive analysis of the interaction between missing data types, imputation methods, and supervised classifiers. *Expert Systems with Applications*, 89:52–65.

[George and Vidyapeetham, 2012] George, A. and Vidyapeetham, A. (2012). Anomaly detection based on machine learning: dimensionality reduction using PCA and classification using SVM. *International Journal of Computer Applications*, 47(21):5–8.

[Ghosh et al., 2017] Ghosh, A., Kumar, H., and Sastry, P. (2017). Robust loss functions under label noise for deep neural networks. In *Proceedings of the AAAI Conference on Artificial Intelligence*.

[Goddard, 2017] Goddard, M. (2017). The EU general data protection regulation (GDPR): European regulation that has a global impact. *International Journal of Market Research*, 59(6):703–705.

[Goh et al., 2016] Goh, G., Cotter, A., Gupta, M., and Friedlander, M. (2016). Satisfying real-world goals with dataset constraints. *arXiv preprint arXiv:1606.07558*.

[Goldfarb-Tarrant et al., 2020] Goldfarb-Tarrant, S., Marchant, R., Sanchez, R. M., Pandya, M., and Lopez, A. (2020). Intrinsic bias metrics do not correlate with application bias. *arXiv preprint arXiv:2012.15859*.

[Gómez-Ríos et al., 2017] Gómez-Ríos, A., Luengo, J., and Herrera, F. (2017). A study on the noise label influence in boosting algorithms: AdaBoost, GBM and XGBoost. In *International Conference on Hybrid Artificial Intelligence Systems*, pages 268–280. Springer.

[Gretton et al., 2009] Gretton, A., Smola, A., Huang, J., Schmittfull, M., Borgwardt, K., and Schölkopf, B. (2009). Covariate shift by kernel mean matching. *Dataset Shift in Machine Learning*, 3(4):5.

[Groenwold et al., 2012] Groenwold, R. H., White, I. R., Donders, A. R. T., Carpenter, J. R., Altman, D. G., and Moons, K. G. (2012). Missing covariate data in clinical research: when and when not to use the missing-indicator method for analysis. *Canadian Medical Association Journal*, 184(11):1265–1269.

[Guo et al., 2017] Guo, C., Pleiss, G., Sun, Y., and Weinberger, K. Q. (2017). On calibration of modern neural networks. In *International Conference on Machine Learning*, pages 1321–1330. PMLR.

[Gupta et al., 2016] Gupta, A., Gusain, K., and Popli, B. (2016). Verifying the value and veracity of extreme gradient boosted decision trees on a variety of datasets. In *2016 11th International Conference on Industrial and Information Systems (ICIIS)*, pages 457–462. Ieee.

[Hajian et al., 2015] Hajian, S., Domingo-Ferrer, J., Monreale, A., Pedreschi, D., and Giannotti, F. (2015). Discrimination-and privacy-aware patterns. *Data Mining and Knowledge Discovery*, 29(6):1733–1782.

[Hanley and McNeil, 1982] Hanley, J. A. and McNeil, B. J. (1982). The meaning and use of the area under a receiver operating characteristic (ROC) curve. *Radiology*, 143(1):29–36.

[Hardt et al., 2016] Hardt, M., Price, E., and Srebro, N. (2016). Equality of opportunity in supervised learning. *arXiv preprint arXiv:1610.02413*.

[Hastie et al., 2009] Hastie, T., Tibshirani, R., and Friedman, J. (2009). Boosting and additive trees. In *The Elements of Statistical Learning*, pages 337–387. Springer.

[Heskes, 1994] Heskes, T. (1994). *The use of being stubborn and introspective*, pages 1184–1200. Citeseer.

[Holzinger et al., 2017] Holzinger, A., Biemann, C., Pattichis, C. S., and Kell, D. B. (2017). What do we need to build explainable AI systems for the medical domain? *arXiv preprint arXiv:1712.09923*.

[Huberman and Langholz, 1999] Huberman, M. and Langholz, B. (1999). Application of the missing-indicator method in matched case-control studies with incomplete data. *American Journal of Epidemiology*, 150(12):1340–1345.

[Jacobusse and Veenman, 2016] Jacobusse, G. and Veenman, C. (2016). On selection bias with imbalanced classes. In *International Conference on Discovery Science*, pages 325–340. Springer.

[James et al., 2013] James, G., Witten, D., Hastie, T., and Tibshirani, R. (2013). *An Introduction to Statistical Learning*, volume 112. Springer.

[Jeatrakul et al., 2010] Jeatrakul, P., Wong, K. W., and Fung, C. C. (2010). Data cleaning for classification using misclassification analysis. *Journal of Advanced Computational Intelligence and Intelligent Informatics*, 14(3):297–302.

[Jiang et al., 2020] Jiang, R., Pacchiano, A., Stepleton, T., Jiang, H., and Chiappa, S. (2020). Wasserstein fair classification. In *Uncertainty in Artificial Intelligence*, pages 862–872. PMLR.

[John, 1995] John, G. H. (1995). Robust decision trees: removing outliers from databases. In *KDD*, volume 95, pages 174–179.

[Jones, 1979] Jones, D. S. (1979). *Elementary Information Theory*. Clarendon Press.

[Kamiran et al., 2010] Kamiran, F., Calders, T., and Pechenizkiy, M. (2010). Discrimination aware decision tree learning. In *2010 IEEE International Conference on Data Mining*, pages 869–874. IEEE.

[Kamishima et al., 2012] Kamishima, T., Akaho, S., Asoh, H., and Sakuma, J. (2012). Fairness-aware classifier with prejudice remover regularizer. In *Joint European Conference on Machine Learning and Knowledge Discovery in Databases*, pages 35–50. Springer.

[Kirch, 2008] Kirch, W. (2008). Pearson's correlation coefficient. *Encyclopedia of Public Health*, pages 1090–1091.

[Kleinbaum and Klein, 2010] Kleinbaum, D. G. and Klein, M. (2010). Modeling strategy guidelines. In *Logistic Regression*, pages 165–202. Springer.

[Kleinberg et al., 2016] Kleinberg, J., Mullainathan, S., and Raghavan, M. (2016). Inherent trade-offs in the fair determination of risk scores. *arXiv preprint arXiv:1609.05807*.

[Knol et al., 2010] Knol, M. J., Janssen, K. J., Donders, A. R. T., Egberts, A. C., Heerdink, E. R., Grobbee, D. E., Moons, K. G., and Geerlings, M. I. (2010).

Unpredictable bias when using the missing indicator method or complete case analysis for missing confounder values: an empirical example. *Journal of Clinical Epidemiology*, 63(7):728–736.

[Kohavi et al., 1996] Kohavi, R., Wolpert, D. H., et al. (1996). Bias plus variance decomposition for zero-one loss functions. In *ICML*, volume 96, pages 275–83.

[Koplowitz and Brown, 1981] Koplowitz, J. and Brown, T. A. (1981). On the relation of performance to editing in nearest neighbor rules. *Pattern Recognition*, 13(3):251–255.

[Kylberg and Sintorn, 2013] Kylberg, G. and Sintorn, I.-M. (2013). Evaluation of noise robustness for local binary pattern descriptors in texture classification. *EURASIP Journal on Image and Video Processing*, 2013(1):1–20.

[Lee, 2019] Lee, J.-S. (2019). AUC4.5: AUC-based C4.5 decision tree algorithm for imbalanced data classification. *IEEE Access*, 7:106034–106042.

[Li et al., 2018] Li, R., Zhang, Y., Tuo, Y., and Chang, P. (2018). A novel method for detecting telecom fraud user. In *2018 3rd International Conference on Information Systems Engineering (ICISE)*, pages 46–50. IEEE.

[Li et al., 2015] Li, W., Mo, W., Zhang, X., Squiers, J. J., Lu, Y., Sellke, E. W., Fan, W., DiMaio, J. M., and Thatcher, J. E. (2015). Outlier detection and removal improves accuracy of machine learning approach to multispectral burn diagnostic imaging. *Journal of Biomedical Optics*, 20(12):121305.

[Lipton et al., 2016] Lipton, Z. C., Kale, D., and Wetzel, R. (2016). Directly modeling missing data in sequences with rnns: Improved classification of clinical time series. In *Machine learning for healthcare conference*, pages 253–270. PMLR.

[Little, 1988] Little, R. J. (1988). A test of missing completely at random for multivariate data with missing values. *Journal of the American Statistical Association*, 83(404):1198–1202.

[Little and Rubin, 2019] Little, R. J. and Rubin, D. B. (2019). *Statistical Analysis with Missing Data*, volume 793. John Wiley & Sons.

[Little et al., 2014] Little, T. D., Jorgensen, T. D., Lang, K. M., and Moore, E. W. G. (2014). On the joys of missing data. *Journal of Pediatric Psychology*, 39(2):151–162.

[Liu et al., 2008] Liu, F. T., Ting, K. M., and Zhou, Z.-H. (2008). Isolation forest. In *2008 8th IEEE International Conference on Data Mining*, pages 413–422. IEEE.

[Liu et al., 2012] Liu, F. T., Ting, K. M., and Zhou, Z.-H. (2012). Isolation-based anomaly detection. *ACM Transactions on Knowledge Discovery from Data (TKDD)*, 6(1):1–39.

[Liu et al., 2017] Liu, H., Li, X., Li, J., and Zhang, S. (2017). Efficient outlier detection for high-dimensional data. *IEEE Transactions on Systems, Man, and Cybernetics: Systems*, 48(12):2451–2461.

[Liu et al., 2016] Liu, J., Li, J., Li, W., and Wu, J. (2016). Rethinking big data: a review on the data quality and usage issues. *ISPRS Journal of Photogrammetry and Remote Sensing*, 115:134–142.

[Liu and Özsu, 2009] Liu, L. and Özsu, M. T. (2009). *Encyclopedia of Database Systems*, volume 6. Springer New York, NY, USA.

[Liu and Tao, 2015] Liu, T. and Tao, D. (2015). Classification with noisy labels by importance reweighting. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 38(3):447–461.

[Liu et al., 2014] Liu, W., Hua, G., and Smith, J. R. (2014). Unsupervised one-class learning for automatic outlier removal. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 3826–3833.

[Malla et al., 2018] Malla, L., Perera-Salazar, R., McFadden, E., Ogero, M., Stepniewska, K., and English, M. (2018). Handling missing data in propensity score estimation in comparative effectiveness evaluations: a systematic review. *Journal of Comparative Effectiveness Research*, 7(3):271–279.

[Manwani and Sastry, 2013] Manwani, N. and Sastry, P. (2013). Noise tolerance under risk minimization. *IEEE Transactions on Cybernetics*, 43(3):1146–1151.

[Mariet et al., 2016] Mariet, Z., Harding, R., Madden, S., et al. (2016). Outlier detection in heterogeneous datasets using automatic tuple expansion. Technical report, Massachusetts Institute of Technology.

[Martínez-Plumed et al., 2019] Martínez-Plumed, F., Contreras-Ochando, L., Ferri, C., Orallo, J. H., Kull, M., Lachiche, N., Quintana, M. J. R., and Flach, P. A. (2019). CRISP-DM twenty years later: From data mining processes to data science trajectories. *IEEE Transactions on Knowledge and Data Engineering*.

[Meertens et al., 2021] Meertens, Q. A. et al. (2021). *Misclassification bias in statistical learning*. PhD thesis, Leiden University.

[Miranda et al., 2009] Miranda, A. L., Garcia, L. P. F., Carvalho, A. C., and Lorena, A. C. (2009). Use of classification algorithms in noise detection and elimination. In *International Conference on Hybrid Artificial Intelligence Systems*, pages 417–424. Springer.

[Mitchell, 1997] Mitchell, T. (1997). *Machine Learning*. New York: McGraw Hill.

[Mok et al., 2010] Mok, M. S., Sohn, S. Y., and Ju, Y. H. (2010). Random effects logistic regression model for anomaly detection. *Expert Systems with Applications*, 37(10):7162–7166.

[Müller and Markert, 2019] Müller, N. M. and Markert, K. (2019). Identifying mislabeled instances in classification datasets. In *2019 International Joint Conference on Neural Networks (IJCNN)*, pages 1–8. IEEE.

[Mulongo et al., 2020] Mulongo, J., Atemkeng, M., Ansah-Narh, T., Rockefeller, R., Nguegnang, G. M., and Garuti, M. A. (2020). Anomaly detection in power generation plants using machine learning and neural networks. *Applied Artificial Intelligence*, 34(1):64–79.

[Muniyandi et al., 2012] Muniyandi, A. P., Rajeswari, R., and Rajaram, R. (2012). Network anomaly detection by cascading k-means clustering and C4.5 decision tree algorithm. *Procedia Engineering*, 30:174–182.

[Narkhede, 2018] Narkhede, S. (2018). Understanding AUC-ROC curve. *Towards Data Science*, 26:220–227.

[Naseer et al., 2018] Naseer, S., Saleem, Y., Khalid, S., Bashir, M. K., Han, J., Iqbal, M. M., and Han, K. (2018). Enhanced network anomaly detection based on deep neural networks. *IEEE Access*, 6:48231–48246.

[Niculescu-Mizil and Caruana, 2005] Niculescu-Mizil, A. and Caruana, R. (2005). Predicting good probabilities with supervised learning. In *Proceedings of the 22nd International Conference on Machine Learning*, pages 625–632.

[Oliphant, 2007] Oliphant, T. E. (2007). Python for scientific computing. *Computing in Science & Engineering*, 9(3):10–20.

[Olson, 2003] Olson, J. E. (2003). *Data Quality: the Accuracy Dimension*. Elsevier.

[Pafka, 2019] Pafka, S. (2019). Benchm-ml. `https://github.com/szilard/benchm-ml`.

[Painsky and Wornell, 2018] Painsky, A. and Wornell, G. (2018). On the universality of the logistic loss function. In *2018 IEEE International Symposium on Information Theory (ISIT)*, pages 936–940. IEEE.

[Paris MoU, 2020] Paris MoU (2020). Current flag performance list. https://www.parismou.org/detentions-banning/white-grey-and-black-list.

[Park and Kwak, 2016] Park, S. and Kwak, N. (2016). Analysis on the dropout effect in convolutional neural networks. In *Asian Conference on Computer Vision*, pages 189–204. Springer.

[Pechenizkiy et al., 2006] Pechenizkiy, M., Tsymbal, A., Puuronen, S., and Pechenizkiy, O. (2006). Class noise and supervised learning in medical domains: the effect of feature extraction. In *19th IEEE Symposium on Computer-based Medical Systems (CBMS'06)*, pages 708–713. IEEE.

[Pedersen et al., 2017] Pedersen, A. B., Mikkelsen, E. M., Cronin-Fenton, D., Kristensen, N. R., Pham, T. M., Pedersen, L., and Petersen, I. (2017). Missing data and multiple imputation in clinical epidemiological research. *Clinical Epidemiology*, 9:157.

[Pedregosa et al., 2011] Pedregosa, F., Varoquaux, G., Gramfort, A., Michel, V., Thirion, B., Grisel, O., Blondel, M., Prettenhofer, P., Weiss, R., Dubourg, V., et al. (2011). Scikit-learn: machine learning in Python. *The Journal of Machine Learning Research*, 12:2825–2830.

[Pereira Barata, 2020] Pereira Barata, A. (2020). Crosslier detection. https://github.com/pereirabarataap/crosslier-detection.

[Pereira Barata, 2021] Pereira Barata, A. (2021). Fair tree classifier. https://github.com/pereirabarataap/fair_tree_classifier.

[Pereira Barata et al., 2018a] Pereira Barata, A., de Bruin, G. J., Takes, F. W., Veenman, C. J., and van den Herik, H. J. (2018a). Data-driven risk assessment in infrastructure networks. In *ICT.open*.

[Pereira Barata et al., 2018b] Pereira Barata, A., de Bruin, G. J., Takes, F. W., Veenman, C. J., and van den Herik, H. J. (2018b). Finding anomalies in waste transportation data with supervised category models. In *2018 27th Belgian Dutch Conference on Machine Learning (BeNeLearn)*.

[Pereira Barata et al., 2019] Pereira Barata, A., Takes, F. W., van den Herik, H. J., and Veenman, C. J. (2019). Imputation methods outperform missing-indicator for data missing completely at random. In *2019 International Conference on Data Mining Workshops (ICDMW)*, pages 407–414. IEEE.

[Pereira Barata et al., 2021] Pereira Barata, A., Takes, F. W., van den Herik, H. J., and Veenman, C. J. (2021). The eXPose approach to crosslier detection. In

2020 25th International Conference on Pattern Recognition (ICPR), pages 2312–2319. IEEE.

[Pessach and Shmueli, 2020] Pessach, D. and Shmueli, E. (2020). Algorithmic fairness. *arXiv preprint arXiv:2001.09784*.

[Pessach and Shmueli, 2022] Pessach, D. and Shmueli, E. (2022). A review on fairness in machine learning. *ACM Computing Surveys (CSUR)*, 55(3):1–44.

[Platt et al., 1999] Platt, J. et al. (1999). Probabilistic outputs for support vector machines and comparisons to regularized likelihood methods. *Advances in Large Margin Classifiers*, 10(3):61–74.

[Port of Rotterdam Authority, 2021] Port of Rotterdam Authority (2021). Facts and figures. `https://www.portofrotterdam.com/sites/default/files/2021-06/facts-and-figures-port-of-rotterdam.pdf`.

[Quy et al., 2021] Quy, T. L., Roy, A., Iosifidis, V., and Ntoutsi, E. (2021). A survey on datasets for fairness-aware machine learning. *arXiv preprint arXiv:2110.00530*.

[Rausand, 2013] Rausand, M. (2013). *Risk assessment: theory, methods, and applications*, volume 115. John Wiley & Sons.

[Rekatsinas et al., 2017] Rekatsinas, T., Chu, X., Ilyas, I. F., and Ré, C. (2017). Holoclean: holistic data repairs with probabilistic inference. *arXiv preprint arXiv:1702.00820*.

[Ren et al., 2018] Ren, M., Zeng, W., Yang, B., and Urtasun, R. (2018). Learning to reweight examples for robust deep learning. In *International Conference on Machine Learning*, pages 4334–4343. PMLR.

[Richardson, 2022] Richardson, A. (2022). Biased data lead to biased algorithms. *CMAJ*, 194(9):E341–E341.

[Rieger et al., 2010] Rieger, A., Hothorn, T., and Strobl, C. (2010). Random forests with missing values in the covariates. Technical Report 79, University of Munich.

[Robertson, 2008] Robertson, S. (2008). A new interpretation of average precision. In *Proceedings of the 31st Annual International ACM SIGIR Conference on Research and Development in Information Retrieval*, pages 689–690.

[Rosario, 2004] Rosario, D. S. (2004). Highly effective logistic regression model for signal (anomaly) detection. In *2004 IEEE International Conference on Acoustics, Speech, and Signal Processing*, volume 5, pages V–817. IEEE.

[Rose and Fischer, 2011] Rose, L. T. and Fischer, K. W. (2011). Garbage in, garbage out: having useful data is everything. *Measurement: Interdisciplinary Research & Perspective*, 9(4):222–226.

[Sáez et al., 2014] Sáez, J. A., Galar, M., Luengo, J., and Herrera, F. (2014). Analyzing the presence of noise in multi-class problems: alleviating its influence with the one-vs-one decomposition. *Knowledge and Information Systems*, 38(1):179–206.

[Saito and Rehmsmeier, 2015] Saito, T. and Rehmsmeier, M. (2015). The precision-recall plot is more informative than the ROC plot when evaluating binary classifiers on imbalanced datasets. *PLOS ONE*, 10(3):e0118432.

[Samuel, 1959] Samuel, A. L. (1959). Some studies in machine learning using the game of checkers. *IBM Journal of research and development*, 3(3):210–229.

[Santos et al., 2019a] Santos, J. D., Chebotarov, D., McNally, K. L., Bartholomé, J., Droc, G., Billot, C., and Glaszmann, J. C. (2019a). Fine scale genomic signals of admixture and alien introgression among Asian rice landraces. *Genome Biology and Evolution*, 11(5):1358–1373.

[Santos et al., 2019b] Santos, M. S., Pereira, R. C., Costa, A. F., Soares, J. P., Santos, J., and Abreu, P. H. (2019b). Generating synthetic missing data: a review by missing mechanism. *IEEE Access*, 7:11651–11667.

[Sarker, 2021] Sarker, I. H. (2021). Machine learning: algorithms, real-world applications and research directions. *SN Computer Science*, 2(3):1–21.

[Schlomer et al., 2010] Schlomer, G. L., Bauman, S., and Card, N. A. (2010). Best practices for missing data management in counseling psychology. *Journal of Counseling Psychology*, 57(1):1.

[Schrider and Kern, 2018] Schrider, D. R. and Kern, A. D. (2018). Supervised machine learning for population genetics: a new paradigm. *Trends in Genetics*, 34(4):301–312.

[Scott, 2015] Scott, C. (2015). A rate of convergence for mixture proportion estimation, with application to learning from noisy labels. In *Artificial Intelligence and Statistics*, pages 838–846. PMLR.

[Segata et al., 2010] Segata, N., Blanzieri, E., Delany, S. J., and Cunningham, P. (2010). Noise reduction for instance-based learning with a local maximal margin approach. *Journal of Intelligent Information Systems*, 35(2):301–331.

[Sharma et al., 2017] Sharma, N., Verlekar, P., Ashary, R., and Zhiquan, S. (2017). Regularization and feature selection for large dimensional data. *arXiv preprint arXiv:1712.01975*.

[Sidi et al., 2012] Sidi, F., Panahy, P. H. S., Affendey, L. S., Jabar, M. A., Ibrahim, H., and Mustapha, A. (2012). Data quality: a survey of data quality dimensions. In *2012 International Conference on Information Retrieval & Knowledge Management*, pages 300–304. IEEE.

[Stone, 1974] Stone, M. (1974). Cross-validatory choice and assessment of statistical predictions. *Journal of the Royal Statistical Society: Series B (Methodological)*, 36(2):111–133.

[Subudhi and Panigrahi, 2020] Subudhi, S. and Panigrahi, S. (2020). Use of optimized fuzzy c-means clustering and supervised classifiers for automobile insurance fraud detection. *Journal of King Saud University-Computer and Information Sciences*, 32(5):568–575.

[Sun et al., 2007] Sun, J.-w., Zhao, F.-y., Wang, C.-j., and Chen, S.-f. (2007). Identifying and correcting mislabeled training instances. In *Future Generation Communication and Networking (FGCN 2007)*, volume 1, pages 244–250. IEEE.

[Tax and Duin, 2002] Tax, D. M. and Duin, R. P. (2002). Using two-class classifiers for multiclass classification. In *Object Recognition Supported by User Interaction for Service Robots*, volume 2, pages 124–127. IEEE.

[Teng, 2000] Teng, C. M. (2000). Evaluating noise correction. In *Pacific Rim International Conference on Artificial Intelligence*, pages 188–198. Springer.

[Teng, 2001] Teng, C.-M. (2001). A comparison of noise handling techniques. In *FLAIRS Conference*, pages 269–273.

[Thongkam et al., 2008] Thongkam, J., Xu, G., Zhang, Y., and Huang, F. (2008). Support vector machine for outlier detection in breast cancer survivability prediction. In *Asia-Pacific Web Conference*, pages 99–109. Springer.

[Twala, 2009] Twala, B. (2009). An empirical comparison of techniques for handling incomplete data using decision trees. *Applied Artificial Intelligence*, 23(5):373–405.

[Van Buuren, 2018] Van Buuren, S. (2018). *Flexible Imputation of Missing Data*. CRC Press.

[Van der Heijden et al., 2006] Van der Heijden, G. J., Donders, A. R. T., Stijnen, T., and Moons, K. G. (2006). Imputation of missing values is superior to complete case analysis and the missing-indicator method in multivariable diagnostic research: a clinical example. *Journal of Clinical eEidemiology*, 59(10):1102–1109.

[Van Rijn and Hutter, 2018] Van Rijn, J. N. and Hutter, F. (2018). Hyperparameter importance across datasets. In *Proceedings of the 24th ACM SIGKDD International Conference on Knowledge Discovery & Data Mining*, pages 2367–2376.

[Vanschoren et al., 2014] Vanschoren, J., Van Rijn, J. N., Bischl, B., and Torgo, L. (2014). OpenML: networked science in machine learning. *ACM SIGKDD Explorations Newsletter*, 15(2):49–60.

[Vapnik, 2013] Vapnik, V. (2013). *The Nature of Statistical Learning Theory*. Springer Science & Business Media.

[Venkatasubramanian, 2019] Venkatasubramanian, S. (2019). Algorithmic fairness: measures, methods and representations. In *Proceedings of the 38th ACM SIGMOD-SIGACT-SIGAI Symposium on Principles of Database Systems*, pages 481–481.

[Venkatesh and Anuradha, 2019] Venkatesh, B. and Anuradha, J. (2019). A review of feature selection and its methods. *Cybernetics and Information Technologies*, 19(1):3–26.

[Vespe et al., 2012] Vespe, M., Visentini, I., Bryan, K., and Braca, P. (2012). Unsupervised learning of maritime traffic patterns for anomaly detection. In *9th IET Data Fusion & Target Tracking Conference (DF & TT 2012): Algorithms & Applications*. IET.

[Vollmer et al., 2017] Vollmer, M., Sodmann, P., Caanitz, L., Nath, N., and Kaderali, L. (2017). Can supervised learning be used to classify cardiac rhythms? In *2017 Computing in Cardiology (CinC)*, pages 1–4. IEEE.

[Wang et al., 2019] Wang, X., Kodirov, E., Hua, Y., and Robertson, N. M. (2019). Derivative manipulation for general example weighting. *arXiv preprint arXiv:1905.11233*.

[Wang, 2005] Wang, Y. (2005). A multinomial logistic regression modeling approach for anomaly intrusion detection. *Computers & Security*, 24(8):662–674.

[Wilcoxon, 1992] Wilcoxon, F. (1992). Individual comparisons by ranking methods. In *Breakthroughs in Statistics*, pages 196–202. Springer.

[Wilson and Martinez, 2000] Wilson, D. R. and Martinez, T. R. (2000). Reduction techniques for instance-based learning algorithms. *Machine Learning*, 38(3):257–286.

[Wolpert and Macready, 1997] Wolpert, D. H. and Macready, W. G. (1997). No free lunch theorems for optimization. *IEEE Transactions on Evolutionary Computation*, 1(1):67–82.

[Woodworth et al., 2017] Woodworth, B., Gunasekar, S., Ohannessian, M. I., and Srebro, N. (2017). Learning non-discriminatory predictors. In *Conference on Learning Theory*, pages 1920–1953. PMLR.

[Wu, 1995] Wu, X. (1995). *Knowledge Acquisition from Databases*. Intellect books.

[Xia et al., 2017] Xia, Y., Liu, C., Li, Y., and Liu, N. (2017). A boosted decision tree approach using Bayesian hyper-parameter optimization for credit scoring. *Expert Systems with Applications*, 78:225–241.

[Xiang and Min, 2010] Xiang, G. and Min, W. (2010). Applying semi-supervised cluster algorithm for anomaly detection. In *2010 Third International Symposium on Information Processing*, pages 43–45. IEEE.

[Xu et al., 2018] Xu, X., Liu, H., Li, L., and Yao, M. (2018). A comparison of outlier detection techniques for high-dimensional data. *International Journal of Computational Intelligence Systems*, 11(1):652–662.

[Yin and Dong, 2011] Yin, H. and Dong, H. (2011). The problem of noise in classification: past, current and future work. In *2011 IEEE 3rd International Conference on Communication Software and Networks*, pages 412–416. IEEE.

[Zabihi et al., 2017] Zabihi, M., Rad, A. B., Katsaggelos, A. K., Kiranyaz, S., Narkilahti, S., and Gabbouj, M. (2017). Detection of atrial fibrillation in ECG hand-held devices using a random forest classifier. In *2017 Computing in Cardiology (CinC)*, pages 1–4. IEEE.

[Zafar et al., 2017] Zafar, M. B., Valera, I., Rogriguez, M. G., and Gummadi, K. P. (2017). Fairness constraints: mechanisms for fair classification. In *Artificial Intelligence and Statistics*, pages 962–970. PMLR.

[Zanero and Savaresi, 2004] Zanero, S. and Savaresi, S. M. (2004). Unsupervised learning techniques for an intrusion detection system. In *Proceedings of the 2004 ACM Symposium on Applied Computing*, pages 412–419.

[Zhang and Ntoutsi, 2019] Zhang, W. and Ntoutsi, E. (2019). FAHT: an adaptive dairness-aware decision tree classifier. In *Proceedings of the 28th International Joint Conference on Artificial Intelligence, IJCAI-19*, pages 1480–1486.

[Zhang et al., 2006] Zhang, W., Rekaya, R., and Bertrand, K. (2006). A method for predicting disease subtypes in presence of misclassification among training samples using gene expression: application to human breast cancer. *Bioinformatics*, 22(3):317–325.

[Zhang, 2016] Zhang, Z. (2016). Missing data imputation: focusing on single imputation. *Annals of Translational Medicine*, 4(1).

[Zhu and Wu, 2004] Zhu, X. and Wu, X. (2004). Class noise vs. attribute noise: a quantitative study. *Artificial Intelligence Review*, 22(3):177–210.