Handling Veracity of SVM Predictions

Marcelo Loor[1⋆][0000−0002−4027−1628], Ana Tapia-Rosero[2][0000−0002−5326−4750], and Guy De Tré[1][0000−0002−1283−1915]

[1] Ghent University, Dept. of Telecommunications and Information Processing, Sint-Pietersnieuwstraat 41, B-9000, Ghent, Belgium
{Marcelo.Loor,Guy.DeTre}@UGent.be
[2] ESPOL University, Dept. of Electrical and Computer Engineering, Campus Gustavo Galindo V., Km. 30.5 Via Perimetral, Guayaquil, Ecuador
atapia@espol.edu.ec

Abstract. Any concern about the veracity of predictions made by artificial intelligence (AI) systems might deter decision makers from using them to support their decisions. To dismiss such concerns in AI systems based on support vector machines (SVMs), in this paper we explore the use of L-grades for dealing with the veracity of SVM predictions and propose a novel method for obtaining those grades. An illustrative example shows how an L-grade can be used for denoting to what extent an SVM prediction can be trusted, as well as how to obtain L-grades inside a binary classification process.

Keywords: Trustworthy artificial intelligence · Explainable artificial intelligence · Support vector machines · L-grades · Z-numbers.

1 Introduction

Since transparency is deemed a key requirement for *trustworthy artificial intelligence* (TAI) [8, 12], *artificial intelligence* (AI) systems that yield opaque predictions can be banned from situations where decision makers must justify their resolutions [11]. A major problem in this regard is that AI systems based on *support vector machines* (SVMs) [23] can be banned from those situations since their predictions could be difficult for decision makers to understand and, thus, trust [2, 10].

To amend that, several research lines aiming for the improvement of the interpretability of SVM predictions have been followed. One of those lines is in the direction of building simplified twin models with fuzzy rules extracted from SVM models [1, 18]. Another line, which can also be applied to other knowledge models, is directed toward the construction of local models that behave like the original SVM models near the vicinity of the objects under study [9, 19]. A rather new line of research regards the contextualization of SVM predictions by means of contextualized evaluations [13, 14] or through contextualized knowledge models [16].

⋆ Corresponding author.

As a complement to the aforementioned research lines, in this paper we explore how the addition of a veracity component to SVM predictions can help to improve transparency. More specifically, we look into the use of *L-grades* [6] for handling the veracity of SVM predictions and propose a novel method for eliciting such grades. An L-grade, say $l = (s, c) \in [0, 1]^2$, is a special case of a *Z-number* [25] where both components are values of the unit interval $[0, 1]$ which are interpreted as grades that are further processed using fuzzy logic. Analogously to Z-numbers, which have been proposed for dealing with the reliability of (fuzzy) information, herein L-grades are specifically used for coping with the reliability of SVM predictions: while the first component of an L-grade denotes to what extent the evaluation of a proposition having the form "x is (predicted to be) A" is satisfied by x, the second component indicates to what extent the first component can be trusted. For instance, if one is 85% *confident* that the prediction shown in Fig. 1 is satisfied to an extent of 70%, one can use the L-grade $l = (0.70, 0.85)$ to express this information[3].

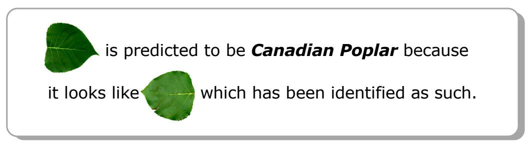

Fig. 1. Example of a prediction with explanation.

Prediction veracity refers to the extent that predictions reflect an accepted knowledge and thus can be trusted (cf. *data veracity* in [3]). As might be noticed in the previous example, L-grades can be used for explicitly expressing the extent to which predictions reflect an accepted knowledge. Thus, L-grades can be used for handling prediction veracity in AI systems based on SVMs. This is a key aspect of our proposal since decision makers can be provided with information that transparently indicates the veracity of a prediction, which can be used along with an explanation like the one shown in Fig. 1 for making an informed decision. It is worth mentioning that providing decision makers with appropriate tools for making such informed decisions is a powerful motivation in our work.

The components of an L-grade, namely the satisfaction and the confidence grades, are deemed truth values and hence lend themselves to logic computation. This is another important aspect of our proposal since the logic framework that has been established for L-grades can be used in situations where a decision needs to be guided by the veracity of an eventual aggregation of two or more SVM predictions. For instance, let p : "x is (predicted to be) A" and q : "x is A" be two propositions related to the predictions made by the contextualized knowledge models P and Q respectively; and let l_p and l_q be the L-grades associated in

[3] The leaves shown in this example are part of the collection described in [24].

that order to p and q. In this case, a decision about x can be guided by a computed veracity of the aggregation of l_p and l_q.

To present our proposal, the paper is structured as follows. A brief summary of the definition and semantics of L-grades is presented in Section 2. In Section 3, we describe how L-grades can be obtained and used for handling prediction veracity in SVM binary classification. An example illustrating this process is presented in Section 4. Related work about veracity in SVM predictions is discussed in Section 5. Finally, we present the conclusions and delineate further work in Section 6.

2 A Brief Summary of L-grades

An L-grade has been conceived as a simplification of a Z-number in which its components are interpreted as grades that can be processed using fuzzy logic [6]. As such, an L-grade is an ordered pair, say $l = (s, c)$, where s and c are considered to be graded values in the unit interval $[0, 1]$.

Semantically, the components of an L-grade are interpreted as follows. The first component, namely s, is called *satisfaction grade* and denotes to what extent an object satisfies a given proposition: while $s = 0$ means that the proposition is not satisfied at all, $s = 1$ means that the proposition is completely satisfied. The second component, namely c, is called *confidence grade* and denotes to what extent the satisfaction grade can be trusted: while $c = 0$ means that (the value of) s cannot be trusted at all, $c = 1$ means that s can be fully trusted. For instance, consider the L-grades $l_P = (1, 1)$ and $l_Q = (1, 0.5)$ given in that order by two persons, say P and Q, after evaluating the proposition "x is A". According to the satisfaction grades in l_p and l_Q, both P and Q consider that the object x fully satisfies the proposition. However, the confidence grades indicate that while P is completely confident that x fully satisfies the proposition, Q is only half confident.

Notice in the previous example that both s and c denote the degrees of truth of two fuzzy propositions: while s indicates the degree of truth of the proposition "x is A", c indicates the degree of truth of the proposition "s is trustable" Hence, s and c can be processed using a logic framework that is truth functional like the one described in [6].

L-grades can be compared to find the proposition that best suits the preferences of a decision maker. For instance, let $l_P = (s_P, c_P)$ and $l_Q = (s_Q, c_Q)$ two L-grades denoting, in that order, the evaluations of the propositions p_P : "x is A" and p_Q : "x is B". If $s_P > s_Q$ and $c_P > c_Q$, a decision maker will prefer p_P instead of p_Q. The same preference will result if s_P and s_Q are very similar and $c_P > c_Q$. However, if $s_P > s_Q$ and $c_P < c_Q$, the decision maker would need additional hints to choose between p_P and p_Q – the interested reader is referred to [6] for possible ways to handle the latter case.

As might be noticed, an L-grade offers the facility to cope with the veracity of a satisfaction grade. This aspect is used in the next section for dealing with the veracity of SVM predictions in a binary classification process.

3 L-grades in Binary SVM Classification

Classification is usually understood as a process in which one or more objects are evaluated in order to determine (or predict) the most suitable class for them. In situations where only two classes have been specified, such a process is deemed *binary classification*.

Concerning the binary classification process performed by an SVM classifier, the idea of a *separating hyperplane* can be used for explaining how the evaluation of the objects is performed. Consider Fig. 2, in which two classes of objects, namely black squares and white circles, as well as three objects, namely x_1, x_2 and x_3, whose classes are unknown are depicted along with a separating hyperplane H. If one is asked to evaluate the extent to which the unknown objects can be classified as black squares, one should say that, since x_3 is very far from H and is located on the side of the black squares, (the proposition) "x_3 is a black square" is fully satisfied. One should also say that, since x_2 is very far from H and is located on the opposite side of the black squares, "x_2 is a black square" is not satisfied. However, one should say that "x_1 is a black square' is hardly satisfied because, although x_1 is located on the side of the black squares, x_1 is very close to H.

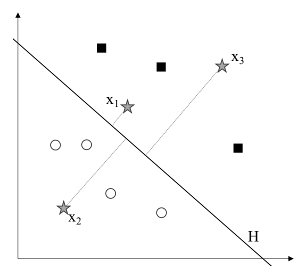

Fig. 2. Idea behind the satisfaction grade binary SVM classification.

As was mentioned in the introduction, we aim to improve the transparency of SVM predictions by the addition of a veracity component, which, as shown in the previous section, can be characterized by an L-grade. Hence, we use the aforementioned idea to obtain the first component of an L-grade, i.e., the satisfaction grade, as follows.

As shown in Fig. 2, the extent to which an object x satisfies the proposition "x is A" is determined by the location of x in relation to the separating hyperplane H. Thus, we compute how far x is from H and its relative position by means of

$$df(\mathbf{x}, \mathbf{w}, b) = \frac{\mathbf{x} \cdot \mathbf{w} + b}{||\mathbf{w}||}, \tag{1}$$

where the normal vector $\mathbf{w}$ and an intercept term b are the components of the hyperplane H, $\mathbf{x}$ is a vector representation of x, '$\cdot$' denotes the inner product, and

$||\mathbf{w}||$ represents the L2 norm of $\mathbf{w}$ – the interested reader is referred to [14] for a detailed explanation about how the components $\mathbf{w}$ and b can be obtained during an SVM learning process. While the magnitude of $df(\mathbf{x}, \mathbf{w}, b)$ is an indication of how far $\mathbf{x}$ is from H, the sign of $df(\mathbf{x}, \mathbf{w}, b)$ is an indication of the side $\mathbf{x}$ is located: a positive sign means that $\mathbf{x}$ is located on the side of (the class) A (specified in the proposition "x is A"), whereas a negative sign means that $\mathbf{x}$ is located on the opposite side of A.

We use Eq. 1 for computing the satisfaction grade of a proposition having the form "x is (predicted to be) A" through

$$s(\mathbf{x}, \mathbf{w}, b) = \begin{cases} 1, & \text{if } df(\mathbf{x}, \mathbf{w}, b) > M \\ 0, & \text{if } df(\mathbf{x}, \mathbf{w}, b) < -M \\ 0.5 + \frac{df(\mathbf{x}, \mathbf{w}, b)}{2M}, & \text{otherwise.} \end{cases} \tag{2}$$

In this equation, $\mathbf{w}$ and b are the components of the hyperplane representing the knowledge model about A that has been built after performing an SVM learning process, and M represents the maximum absolute value computed by Eq. 1 for the objects constituting the training set X_0 used during that learning process, i.e.,

$$M = \max\{|df(\mathbf{x}, \mathbf{w}, b)| : \forall \mathbf{x} \in X_0\}. \tag{3}$$

Notice that $df(\mathbf{x}, \mathbf{w}, b)$ must be in the interval $[-M, M]$ when x belongs to the training set X_0. However, when x belongs to the test set X, $df(\mathbf{x}, \mathbf{w}, b)$ might be outside of $[-M, M]$. Such a situation is handled through the two first cases stated in Eq. 2.

Regarding the second component of an L-grade, i.e., the confidence grade, its value is determined by how well the separating hyperplane breaks up the two classes. Consider Fig. 3, in which an additional separating hyperplane, namely H^*, has been included. If we are asked to evaluate the extent to which x_3 can be classified as black square in relation to H^*, we should say that, since x_3 is very far from H^* and is located on the side of the black squares, "x_3 is a black square" is fully satisfied as was done before with H. Nevertheless, we notice white circles, to be specific x_4 and x_5, on the side of black squares. Thus, our confidence about saying so should decrease in this case.

Since our confidence decreases when the separation of the two classes is not well done, we can use a metric that indicates how well a separating hyperplane breaks up the two classes in SVM classification. Thus, to compute the confidence grade we use

$$c(\mathbf{w}, b) = m(\mathbf{w}, b), \tag{4}$$

where $m(\mathbf{w}, b)$ represents a metric (e.g., accuracy [21] or recall [22]) that measures the performance of the SVM model characterized by the hyperplane whose components are $\mathbf{w}$ and b.

Putting Eq. 2 and Eq. 4 together, we compute the L-grade of an SVM prediction having the form "x is (predicted to be) A" by means of

$$l(\mathbf{x}, \mathbf{w}, b) = \left(s\left(\mathbf{x}, \mathbf{w}, b\right), c\left(\mathbf{w}, b\right)\right), \tag{5}$$

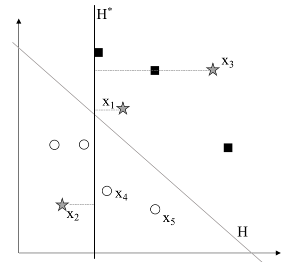

Fig. 3. Idea behind the confidence grade in binary SVM classification.

where, as was mentioned above, $\mathbf{x}$ is a vector representation of x, and $\mathbf{w}$ and b characterize a particular SVM knowledge model about A.

Notice that Eq. 5 can be used for performing a joint evaluation of both the extent s to which an object x satisfies the proposition "x is (predicted to be) A," and the extent c to which this satisfaction grade can be trusted. Such an evaluation can be referred to as an *L-valuation* and can be denoted by the 3-tuple (x, s, c). Moreover, an assembly of such L-valuations can be referred to as *L-information* – cf. the terms *Z-valuation* and *Z-information* introduced by Zadeh in [25]. We consider that L-information extracted from an SVM model can be very useful. For instance, L-information can be used for summarizing the performance of an SVM model and, thus, increasing its transparency.

To make an L-grade more understandable to decision makers, linguistic terms can be used for expressing the values of its components [4, 20]. For instance, the linguistic terms set shown in Fig. 4 can be used for putting into words the satisfaction grades computed with Eq. 2 and, thus, making statements like "the proposition 'x is A' is *enough satisfied* (to an extent of 80%) with a confidence of 95%."

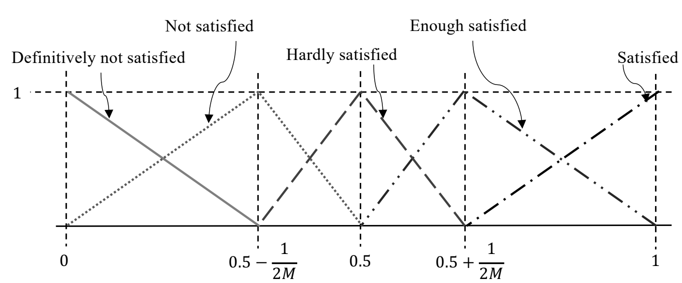

Fig. 4. Linguistic terms associated to satisfaction grades.

In the next section, we illustrate how this method can be used for computing L-grades that denote the veracity of SVM predictions.

4 Illustrative Example

To show how L-grades denote the veracity of SVM predictions in binary classification, a toy data set consisting of 18 black squares and 18 white circles was used. The training set consists of 6 black squares whose coordinates are $(0, 3)$, $(1, 2)$, $(2, 1)$, $(3, 3)$, $(0, 0)$, and $(-1, -1)$, and 5 white circles whose coordinates are $(-1, 0)$, $(-3, 1)$, $(3, -1)$, $(-2, 1)$ and $(0, -4)$. The test set consists of 3 black squares whose coordinates are $(1, 2)$, $(1, 5)$, and $4, 3)$, and 4 white circles whose coordinates are $(-3.5, -3.5)$, $(2, -1)$, $(-1, 1)$, $(-3, 0)$.

The training set and the library presented in [15] were used for training two SVM models with linear kernels and regularization parameters $C_1 = 0.1$ and $C_2 = 1$ respectively. The separating hyperplanes H_1 and H_2 that characterize the resulting models are depicted in Fig. 5. The objects surrounded by a ring are the support vectors.

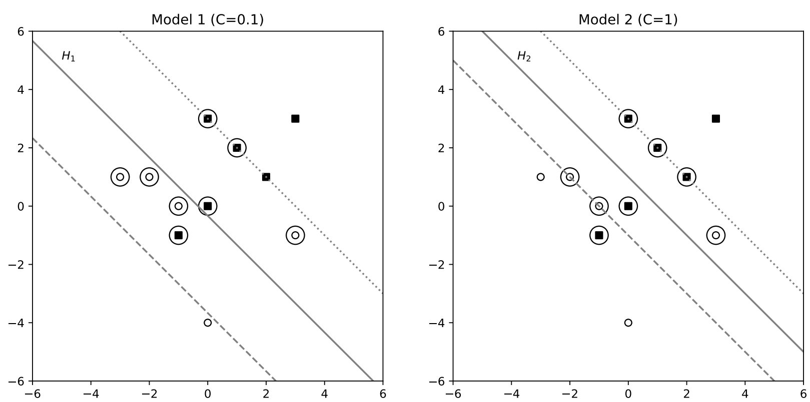

Fig. 5. Separating hyperplanes characterizing the resulting models.

After that, the training set along with the normal vector and the intercept term of each model were used as inputs of Eq. 3 for computing $M_1 = 1.9$ and $M_2 = 2.5$. These values were used as inputs in Eq. 2 for computing the extent to which each object x in the test set satisfies the proposition "x is A", where A represents the black squares.

Then, accuracy was established as the metric that measures the performance of the models and was used along with the test set for computing their confidence grades with Eq. 4. The computed confidence grades for Model 1 and Model 2 are 0.7143 and 0.8571 respectively.

Finally, the computed satisfaction grades and confidence grades were used as inputs of Eq. 5 to assemble the resulting L-grades for each object in the test set. The linguistic terms shown in Fig. 4 were used for making the resulting L-grades more understandable.

The results are shown in Fig. 6. Notice in this figure that, instead of just saying "x_0 is (predicted to be) A" (i.e., a black square), one can say that since "x_0 is A" is enough satisfied with a confidence of 85.71% according to Model 2, x_0 is predicted to be a black square. We consider that this additional information about the veracity of this prediction will help a person to make an informed decision.

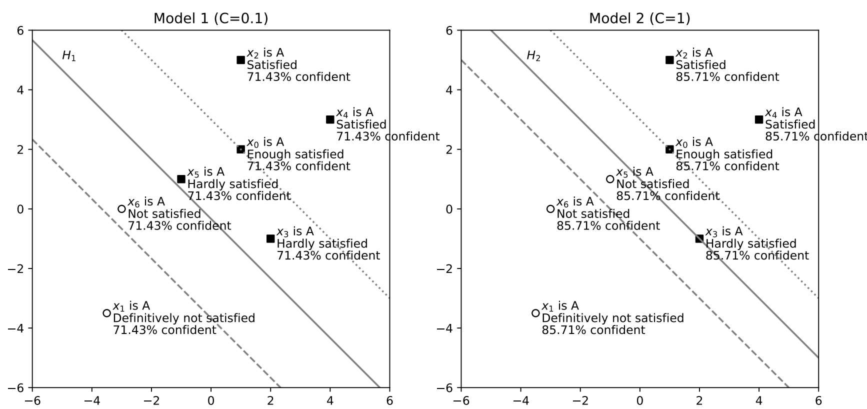

Fig. 6. Computed L-grades.

Notice also in Fig. 6 that the proposition "x_5 is A" is not satisfied with a confidence of 85.71% according to Model 2 and, thus, x_5 is predicted to be a white circle. However, according to Model 1, "x_5 is A" is hardly satisfied with a confidence of 71.43% and, thus, x_5 is predicted to be a black square. In this case, a decision maker will probably accept the prediction made by Model 2 since the confidence grade of this prediction is greater than the confidence grade of the prediction made by Model 1.

As was mentioned in the introduction, an advantage of using L-grades for characterizing the veracity of SVM predictions is that the logic framework established for them can be used. Using that framework, one can, e.g., request the AI system that has been used for building Fig. 6 to show only the predictions that have been enough satisfied with a confidence greater than 80%. As indicated in the previous section, this opens the possibility of extracting useful L-information from SVM models, which is recommended and subject to further study.

5 Related Work

These days it is common to find AI systems that automatically learn from the data, identify patterns, make predictions and some of them can even make predictions from experience without human intervention or assistance, as in the case of AI of things [5]. So, handling explainability and veracity are important tasks that impact decision-making processes, in particular when decision makers are prone to use predictions generated by AI systems. The authors in [26] present a comprehensive review for the use of AI in public governance which includes the benefits of AI to support decision makers.

Regarding the veracity of classification predictions, to the best of our knowledge, the available literature is nonexistent. So, there is a long path to pave regarding veracity in predictions and this proposal contributes towards that direction. However, several articles on data veracity can be found in the literature. One of them focuses on veracity as one of the characteristics of big data [3]. Another describes methods for performing veracity assessments [17]. In [7], the authors present a multi-dimensional framework for handling veracity in multi-criteria decision making.

6 Conclusions

We have studied how the addition of a veracity component into SVM predictions can improve their transparency, which is a fundamental requirement for trustworthy AI. L-grades have been used in this paper for characterizing such a veracity component since those grades enable the assessment of confidence criteria that reflect the performance of SVM knowledge models. Enabling such assessment is deemed crucial for handling veracity of SVM predictions.

We have proposed a novel method for computing L-grades of SVM predictions made during a binary classification process. The satisfaction grade, which is the first component of an L-grade, denotes the extent to which a proposition having the form "x is (predicted to be) A" is satisfied by an object x. Thus, an indication of the location of x in relation to the separating hyperplane that characterizes an SVM model is used for computing the satisfaction grade. The second component, called confidence grade, denotes the extent to which the satisfaction grade is trustable. Hence, a metric that measures the performance of the SVM model used for computing the satisfaction grade can be used in this case for computing the confidence grade.

An example has illustrated how L-grades provide additional information about the veracity of SVM predictions that can help decision makers to make informed decisions. Moreover, the example has shown how linguistic terms can make L-grades easier to understand for them.

The possibility of extracting L-information, which can be obtained and further be processed by means of the logic framework proposed for handling L-grades, has been also exposed in the example. Processing such L-information with the aim of identifying SVM models that are appropriate for a given data set is planned and highly recommended as future work. Studying how such L-information can be used for aggregating two or more contextualized SVM models is also planned.

Acknowledgments. This work was supported by the Flemish Government under the "Onderzoeksprogramma Artificiële Intelligentie (AI) Vlaanderen" programme, and by ESPOL Polytechnic University under the research project "Interpretable Artificial Intelligence (XAI) in Group Decision-Making (FIEC-200-2020)."

References

1. Barakat, N.H., Bradley, A.P.: Rule Extraction from Support Vector Machines: A Sequential Covering Approach. IEEE Transactions on Knowledge and Data Engineering **19**(6), 729–741 (2007). https://doi.org/10.1109/TKDE.2007.190610
2. Barredo Arrieta, A., Díaz-Rodríguez, N., Del Ser, J., Bennetot, A., Tabik, S., Barbado, A., Garcia, S., Gil-Lopez, S., Molina, D., Benjamins, R., Chatila, R., Herrera, F.: Explainable Artificial Intelligence (XAI): Concepts, taxonomies, opportunities and challenges toward responsible AI. Information Fusion **58**, 82–115 (2020). https://doi.org/10.1016/j.inffus.2019.12.012
3. Berti-Equille, L., Ba, M.L.: Veracity of big data: Challenges of cross-modal truth discovery. J. Data and Information Quality **7**(3) (aug 2016). https://doi.org/10.1145/2935753
4. Bordogna, G., Pasi, G.: A fuzzy linguistic approach generalizing boolean information retrieval: A model and its evaluation. Journal of the American Society for Information Science **44**(2), 70–82 (1993)
5. Chang, Z., Liu, S., Xiong, X., Cai, Z., Tu, G.: A survey of recent advances in edge-computing-powered artificial intelligence of things. IEEE Internet of Things Journal **8**(18), 13849–13875 (2021)
6. De Tré, G., Peelman, M., Dujmović, J.: Logic reasoning under data veracity concerns. International Journal of Approximate Reasoning **161**, 108977 (2023). https://doi.org/10.1016/j.ijar.2023.108977
7. De Tré, G., De Mol, R., Bronselaer, A.: Handling veracity in multi-criteria decision-making: A multi-dimensional approach. Information Sciences **460-461**, 541–554 (2018). https://doi.org/10.1016/j.ins.2017.09.008
8. Díaz-Rodríguez, N., Del Ser, J., Coeckelbergh, M., López de Prado, M., Herrera-Viedma, E., Herrera, F.: Connecting the dots in trustworthy Artificial Intelligence: From AI principles, ethics, and key requirements to responsible AI systems and regulation. Information Fusion **99**, 101896 (2023). https://doi.org/10.1016/j.inffus.2023.101896
9. Guidotti, R., Monreale, A., Giannotti, F., Pedreschi, D., Ruggieri, S., Turini, F.: Factual and counterfactual explanations for black box decision making. IEEE Intelligent Systems **34**(6), 14–23 (Nov 2019). https://doi.org/10.1109/MIS.2019.2957223
10. Guidotti, R., Monreale, A., Ruggieri, S., Turini, F., Giannotti, F., Pedreschi, D.: A survey of methods for explaining black box models. ACM Comput. Surv. **51**(5) (Aug 2018). https://doi.org/10.1145/3236009
11. Kaminski, M.E.: The right to explanation, explained. Berkeley Tech. LJ **34**, 189 (2019). https://doi.org/10.15779/Z38TD9N83H

12. Li, B., Qi, P., Liu, B., Di, S., Liu, J., Pei, J., Yi, J., Zhou, B.: Trustworthy AI: From Principles to Practices. ACM Comput. Surv. **55**(9) (jan 2023). https://doi.org/10.1145/3555803
13. Loor, M., De Tré, G.: Explaining Computer Predictions with Augmented Appraisal Degrees. In: 2019 Conference of the International Fuzzy Systems Association and the European Society for Fuzzy Logic and Technology (EUSFLAT 2019). Atlantis Press (2019/08). https://doi.org/10.2991/eusflat-19.2019.24
14. Loor, M., De Tré, G.: Contextualizing support vector machine predictions. International Journal of Computational Intelligence Systems **13**, 1483–1497 (2020). https://doi.org/10.2991/ijcis.d.200910.002
15. Loor, M., Tapia-Rosero, A., De Tré, G.: An Open-Source Software Library for Explainable Support Vector Machine Classification. In: 2022 IEEE International Conference on Fuzzy Systems (FUZZ-IEEE). pp. 1–7 (2022). https://doi.org/10.1109/FUZZ-IEEE55066.2022.9882731
16. Loor, M., Tapia-Rosero, A., De Tré, G.: Contextual Boosting to Explainable SVM Classification. In: Fuzzy Logic and Technology, and Aggregation Operators: 13th Conference of the European Society for Fuzzy Logic and Technology, EUSFLAT 2023, and 12th International Summer School on Aggregation Operators, AGOP 2023, Palma de Mallorca, Spain, September 4–8, 2023, Proceedings. p. 480–491. Springer-Verlag, Berlin, Heidelberg (2023). https://doi.org/10.1007/978-3-031-39965-7_40
17. Lozano, M.G., Brynielsson, J., Franke, U., Rosell, M., Tjörnhammar, E., Varga, S., Vlassov, V.: Veracity assessment of online data. Decision Support Systems **129**, 113132 (2020). https://doi.org/10.1016/j.dss.2019.113132
18. Martens, D., Baesens, B., Van Gestel, T.: Decompositional rule extraction from support vector machines by active learning. IEEE Transactions on Knowledge and Data Engineering **21**(2), 178–191 (2009). https://doi.org/10.1109/TKDE.2008.131
19. Ribeiro, M.T., Singh, S., Guestrin, C.: "Why Should I Trust You?": Explaining the Predictions of Any Classifier. In: Proceedings of the 22nd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining. pp. 1135–1144. KDD '16, ACM, New York, NY, USA (2016). https://doi.org/10.1145/2939672.2939778
20. Rodriguez, R.M., Martinez, L., Herrera, F.: Hesitant fuzzy linguistic term sets for decision making. IEEE Transactions on fuzzy systems **20**(1), 109–119 (2011)
21. Sammut, C., Webb, G.I. (eds.): Accuracy, pp. 9–10. Springer US, Boston, MA (2010). https://doi.org/10.1007/978-0-387-30164-8_3
22. Sammut, C., Webb, G.I. (eds.): Precision and Recall, pp. 781–781. Springer US, Boston, MA (2010). https://doi.org/10.1007/978-0-387-30164-8_652
23. Vapnik, V.N., Vapnik, V.: Statistical learning theory, vol. 1. Wiley New York (1998)
24. Wu, S.G., Bao, F.S., Xu, E.Y., Wang, Y.X., Chang, Y.F., Xiang, Q.L.: A leaf recognition algorithm for plant classification using probabilistic neural network. In: 2007 IEEE International Symposium on Signal Processing and Information Technology. pp. 11–16 (2007). https://doi.org/10.1109/ISSPIT.2007.4458016
25. Zadeh, L.A.: A Note on Z-numbers. Information Sciences **181**(14), 2923–2932 (2011). https://doi.org/10.1016/j.ins.2011.02.022
26. Zuiderwijk, A., Chen, Y.C., Salem, F.: Implications of the use of artificial intelligence in public governance: A systematic literature review and a research agenda. Government Information Quarterly **38**(3), 101577 (2021)