

Towards Explainable Linguistic Summaries

Citation for published version (APA):

Wrede, C., Winands, M. H. M., Smirnov, E., & Wilbik, A. (2023). Towards Explainable Linguistic Summaries. In 2023 IEEE International Conference on Fuzzy Systems (FUZZ) (pp. 1-6). Article 10309707 The IEEE. https://doi.org/10.1109/FUZZ52849.2023.10309707

Document status and date: Published: 17/08/2023

DOI: 10.1109/FUZZ52849.2023.10309707

Document Version: Publisher's PDF, also known as Version of record

Document license: Taverne

Please check the document version of this publication:

 A submitted manuscript is the version of the article upon submission and before peer-review. There can be important differences between the submitted version and the official published version of record. People interested in the research are advised to contact the author for the final version of the publication, or visit the DOI to the publisher's website.

• The final author version and the galley proof are versions of the publication after peer review.

 The final published version features the final layout of the paper including the volume, issue and page numbers.

Link to publication

General rights

Copyright and moral rights for the publications made accessible in the public portal are retained by the authors and/or other copyright owners and it is a condition of accessing publications that users recognise and abide by the legal requirements associated with these riahts.

Users may download and print one copy of any publication from the public portal for the purpose of private study or research.

- You may not further distribute the material or use it for any profit-making activity or commercial gain
 You may freely distribute the URL identifying the publication in the public portal.

If the publication is distributed under the terms of Article 25fa of the Dutch Copyright Act, indicated by the "Taverne" license above, please follow below link for the End User Agreement:

www.umlib.nl/taverne-license

Take down policy

If you believe that this document breaches copyright please contact us at:

repository@maastrichtuniversity.nl

providing details and we will investigate your claim.

Towards Explainable Linguistic Summaries

Carla Wrede*[†], Mark H.M. Winands*, Evgueni Smirnov*, and Anna Wilbik*

* Department of Advanced Computing Sciences (DACS),

Maastricht University, Maastricht, The Netherlands

{c.wrede, m.winands, smirnov, a.wilbik}@maastrichtuniversity.nl

[†] VDL Nedcar B.V., Born, The Netherlands

c.wrede@vdlnedcar.nl

Abstract—As more AI solutions are implemented in every aspect of our lives, the need for Explainable Artificial Intelligence (XAI) rises. Explanations can have different forms, such as a number (or an equation), a figure, or a text. This paper investigates textual explanations to effectively communicate the reason for a decision made by an AI system. In previous works, linguistic summaries, as an example of a textual explanation, have already been tested and shown to have explanatory potential. In this paper, we explore this topic further and present a roadmap for linguistic summaries to become a proper XAI tool as Explainable Linguistic Summaries (XLSs). We discuss the challenges that an XLS has to overcome. We outline possible solutions and state their consequences. We consider different protoforms, the definition of the membership function, heuristics for the selection of XLS and personalizing the explanations as well as options to gain more insights into the explanations.

Index Terms—Explainable Artificial Intelligence, Linguistic Summaries, Roadmap, Explainable Linguistic Summary (XLS)

I. INTRODUCTION

Recent advances in information technology cause companies to discover the value of data, resulting in more and more data being gathered with the hope of analyzing it. The amount of data is beyond human cognitive capabilities and comprehension skills, which can be seen in healthcare data doubling every two years [1]. To process this data, business organizations are using many powerful data mining and knowledge discovery methods, though they still require human understanding. For this purpose, recently there is a big interest in Explainable AI methods [2]. There is an expectation that with those means the users can understand better the machine-made recommendations.

Arrieta et al. [2] distinguish several types of explanations, such as numerical, visual and textual explanations. We focus here on the textual explanations, as they use natural language as the communication means, which is the natural way of communication for humans. We investigate linguistic summaries, as they are able to capture the essence of data (e.g. [3]). Moreover, some experiments showed that some people can make better decisions based on the information presented to them verbally (as a text), instead of as numbers or figures [4]. This is also in line with the results of the OECD numeracy skill study [5], which investigates various skills and

measures the proficiency in numeracy, literacy and problemsolving in adults. This report states that in the worldwide population only 40% reached numerical literacy level 2, being able to understand simple multivariate graphs. The OECD study shows also that people who are struggling with their numeracy have troubles grasping numerical concepts and lack an understanding of more advanced graph types. Therefore, we believe that textual explanations will transfer data and model recommendations to a better understandable format. However, until now, the explanatory capabilities of linguistic summaries were investigated only to a limited extent (cf. [6], [7]). There is a potential for these summaries as ad hoc local explanations but as is the characteristic of natural language, there are many possible ways of creating them.

Therefore, in this paper we want to identify and discuss the challenges that need to be overcome to unlock the full potential of linguistic summaries as good explanations and become Explainable Linguistic Summaries (XLSs). While there are many aspects that have an influence on this, we focus on what we believe are the five core challenges at this moment. Each challenge is described and its selection is motivated alongside potential solutions and possible consequences.

The remainder of this paper is organized as follows. First, in Section II more details about the related work of Explainable Artificial Intelligence are given. Next, section III gives background details about linguistic summaries. Afterwards, Section IV outlines the challenges we see when it comes to designing these XLS to become an XAI tool and addresses how they can realistically be accomplished in practice. Subsequently, Section V contains a discussion of these design challenges and choices with respect to the established related work and possible evaluation metrics. Finally, Section VI concludes this paper and gives an outlook for future research.

II. RELATED WORK

With the increased presence of machine-learning models, there is a need for understanding these models, which has resulted in emerging a new domain of Explainable Artificial Intelligence (XAI) [2], [8], [9]. Generally, the various XAI methods can be divided into two basic categories: model-agnostic XAI methods, which can generate explanations for any type of black-box model and model-specific XAI methods, which are designed for a particular machine-learning model. Moreover, model-agnostic methods can be further

This work has received financial support from the Ministry of Economic Affairs and Climate, under the grant "R&D Mobility Sectors" carried out by the Netherlands Enterprise Agency.

distinguished into local and global methods, depending on whether they provide an explanation for a particular data point or the whole data set.

A number of model-specific XAI methods have already been proposed, e.g., for fuzzy rule-based systems [10], [11], logical formulas [12], counterfactual facts [13], [14], knowledge representation and reasoning [15]–[18], temporal and causal relations in Bayesian networks [19]–[21], and blackbox machine learning algorithms [22], [23]. Regarding modelagnostic approaches, examples are Grad-CAM [24], SHAP [25], LIME [26] or DeepLIFT [27].

In [2] the authors categorize model-agnostic XAI methods in six categories, depending on the explanation means, namely: local explanations, local simplification, feature relevance, explanation by example, visualizations and text explanations. In this paper we use linguistic summaries [28], which we consider as potential local and textual explanations.

III. LINGUISTIC SUMMARIES

We employ the method of a linguistic summary as proposed by Yager [28], which are related to Computing with Words [29]. A linguistic summary is a template-based sentence in semi-natural language. Typically, two protoforms (templates) are used, the simple protoform:

$$Qy$$
's are P (1)

and the extended protoform:

$$Q Ry$$
's are P (2)

where Q is the quantifier, e.g. many, most, P is the summarizer, i.e., a property of the object (or a set of those), Ris the qualifier, i.e., a different object characterization, and y's are the objects to be summarized. The qualifier R is a subset of features and linguistic terms that best describe the summarized data point with the conjunction and. It is also called the subset of feature-linguistic term pairs. Key concepts here are the linguistic variables and the linguistic values. The linguistic variable is based on the linguistic terms describing a certain feature or property of an object, with the linguistic value giving it a certain meaning. For example, a linguistic variable of the feature "age" may have the linguistic terms "new", "pristine", "dated" or "old", each with an underlying meaning as described with a membership function. For a database of cars, a linguistic summary most cars are fast is an example of a simple protoform summary, while most new cars are fast is an example of an extended protoform summary. An example of an extended form with multiple feature-linguistic term pairs is most new and sporty cars are fast.

The basic criterion for evaluating the quality of the linguistic summary is the truth value T, also called the validity of a summary. One possibility to determine its value is to employ Zadeh's calculus [30]. In this case, the truth value for a simple and an extended protoform is calculated respectively as:

$$T(Qy'\text{s are } P) = \mu_Q \left(\frac{1}{n} \left(\sum_{i=1}^n \mu_P(y_i)\right)\right)$$
(3)

$$T(Q \; Ry's \; \text{are} \; P) = \mu_Q \left(\frac{\sum_{i=1}^n \mu_P(y_i) \wedge \mu_R(y_i)}{\sum_{i=1}^n \mu_R(y_i)} \right) \quad (4)$$

where μ . is a membership function of the appropriate linguistic term and n is the number of objects to be summarized. More details about linguistic summarization and different methods for the evaluation of linguistic summaries can be found in [31].

This method has been investigated by many researchers and can summarize different types of structured data: databases [32], [33], time series [34], [35], standardized texts [36], videos [37]–[39], sensor data [40], [3], [41], web logs [42] and event logs [6], [43].

IV. DESIGN CHALLENGES AND CHOICES FOR XLS

The following section outlines the challenges we see for designing linguistic summaries that are automatically generated. We selected five different aspects, that have been deemed as important starting points. Those aspects were generated by looking at developments and challenges of linguistic summaries and discussed by the authors. For each of these points the challenge itself is described, the reason for addressing it and the resulting consequences that follow from it. Additionally, possible solutions to how these challenges could be realistically addressed in practice are given and motivated. A visualization of these challenges can be seen in Figure 1.

A. Protoform

As described in the previous section, linguistic summaries are based on protoforms, also called templates, which give shape to the explanation. Currently, linguistic summaries have only been tested on the tasks of *classification* [7] and *anomaly detection* [6]. For this, the based templates as described by (1) and (2) have been adapted to account for the task-specific needs for such an explanation. The explanatory power has been found to be sufficient, and we assume that linguistic summaries can also be efficiently used on other tasks. For this, it is important to be able to easily adapt the base templates to different tasks, also to establish a recognition value of this type of explanation. While the base templates might already offer a valid explanation, making use of the possibly inherent structure of the different tasks the explanations are used for can streamline the explanations even more. A general similarity between templates means a reduced adaptation for recognizing and understanding the general concept of these types of explanations for the user. Moreover, any gained structural insights from improved tasks-specific explanations could in turn also be more easily transferred to explanations of other tasks.

In the following, we propose exemplary XLSs for four basic tasks of data mining. Please note that these adaptations are specifications of the extended protoform (2) by putting certain conditions on both qualifiers and summarizer, and can also be adapted to more advanced tasks.

1) Classification: In the case of classification, we want to explain why a certain label \hat{C} was assigned by a classification model. Therefore, the template should contain this assigned



Fig. 1. Visualization of the challenges of Explainable Linguistic Summaries.

label, either as the qualifier or as the summarizer C. This would give us two different templates:

$$Q Ry$$
's are \hat{C} (5)

$$Q \hat{C} \text{ are } P$$
 (6)

So far, only (5) has been tested in [7]. The choice was made as it seems more intuitive, and it was perceived well in an initial assessment. The alternative protoform of (6) has yet to be evaluated.

2) Clustering: The task of clustering remains an interesting field of study. In our view, there are two options for building a template for cluster description. The first one puts the assigned cluster C_i in focus, while the second one highlights the shared summarizer of the summarized objects within the same cluster:

$$Q Ry$$
's are in cluster C_i (7)

$$Qy$$
's in C_i are P (8)

3) Anomaly detection: For the case of anomaly detection, as investigated in [6], it is important to find the biggest difference between the anomaly object and the normal object. Therefore, the template should point out the found contrast clearly.

$$Qy$$
's are P and anomaly is not P . (9)

4) Association rules: Finally, linguistic summaries can be used for improving the interpretability of association rules. Both templates (1) and (2) can be employed, however, in a slightly modified fashion.

$$Q's < verb > P \tag{10}$$

$$Q Ry$$
's $\langle verb \rangle P$ (11)

The placeholder $\langle \text{verb} \rangle$ can be instantiated depending on the application. For example, if we would analyze data from car orders we might find an association rule "Lane Assist and Blind Spot detection \rightarrow Adaptive Cruise Control". In this case, $\langle \text{verb} \rangle$ gets value "have" and one possible linguistic summary can be read as follows: *Most SUVs have Adaptive Cruise Control when also having Lane Assist and Blind Spot detection.*

B. Membership function

The next challenge when automatically creating linguistic summaries lies in establishing the definition of the linguistic variables and their respective linguistic values. Just as in other expert systems, the linguistic values need to be defined in a meaningful way. This can be composed of two tasks, namely determining which linguistic values are used and what their underlying meaning is, which is modeled with a membership function.

Current applications were defining the linguistic variables by questioning the domain experts and potential users, which can potentially be both time-intensive as well as difficult when combining different definitions. Usually, the trapezoidal membership function was used, as recommended by Zadeh [44]. They can be defined by only four parameters and are easy to understand by domain experts. Moreover, the impact of the choice of the membership function is relatively limited [45].

In the future, we see the potential for automatizing this step. The ideas coming from Referring Expression Generation (REG) [46] are a promising direction, however, the current solutions cannot deal with objects described by continuous variables.

C. XLS selection

Another aspect to consider is the actual selection of the presented linguistic summary. In the previous works of [6], [7] multiple explanations have been generated that met a certain threshold level for the truth value T. The final presented explanations were handpicked to evaluate certain characteristics of said explanations. Especially when using the extended base template (2), the number of possible linguistic summaries grows quickly, as the combinations of available feature-linguistic term pairs grows [47]. The manual selection of summaries is neither desired nor efficient, as it is a tedious and time-consuming task. Therefore, we would like to find a (heuristic) approach that can automatically select an appropriate linguistic summary as an explanation. The consequence of having such a selection strategy would be to be able to more easily adapt such a (heuristic) approach based on whatever is needed from the explanation, which is possibly both task- and problem-dependent.

From an initial investigation in [7] we already gained some insights into explanation preferences. Firstly, the truth value has a high importance when it comes to the selected explanation and a higher value is preferred over other aspects of an explanation. Secondly, the strength off the quantifier is also influential, as stronger ones are preferred (i.e. "almost all" preferred to "most" preferred to "many"). However, there is no clear consensus on whether shorter summaries or more detailed ones containing more feature combinations are preferred. This might be accountable to individual needs of such an explanation and will be addressed in the next section.

Through experimentation we can test these insights regarding the preference options for all the mentioned design choices of the explainable linguistic summaries. Moreover, multicriteria or multi-attribute decision-making methods could help us rank the interesting summaries [48]. These methods are particularly suited, as they aim to help find a decision when there are conflicting needs present.

Another approach in the same direction is to employ the concept of version spaces [49]. A version space VS in our context is the set of all the linguistic summaries with truth value T greater than a user-defined threshold. As we have already discussed, the number of those linguistic summaries can be big; i.e. the size of VS can be prohibitively large. However, since the space of all possible linguistic summaries is a partially ordered set, we can represent the version space VS with its minimal boundary set and maximal boundary set. In this way, if a user u needs to know all the most specific (general) linguistic summaries, the summaries from

the minimal (maximal) boundary set of version space VS can be employed. If the user needs to know all the plausible linguistic summaries from VS, the summaries from both boundary sets of VS can be employed (since any plausible candidate is bounded by a pair of the most specific and the most general linguistic summaries).

D. Personalizing

While it is already effective to have a general selection strategy to choose a specific linguistic summary as an explanation, it is difficult to evaluate how it is perceived by the end user receiving it. As each person is different, therefore the user perception can become quite subjective. One can distinguish between three different groups, novice, domain experts, and AI experts, as proposed by [50]. Whereas a certain explanation might be sufficient for a floor worker, it may leave many open questions for an office worker. Therefore, we propose to put the end user receiving the explanation in the focus.

Personalizing through Explainable Linguistic Summaries can be realized in two ways, either during the membership function definition or when selecting an XLS. To personalize the membership function, two commonly used approaches are the use of interval fuzzy sets and type 2 fuzzy sets. Type 2 fuzzy sets enable to include uncertainty in the value of membership degree [51]. In the work of [52] type 2 linguistic summaries have already been proposed.

Personalizing for XLS selection depends on the user type, and is essentially a problem of selecting the most plausible linguistic summaries for a concrete user u (type). We propose to solve this problem by learning a version space VSu for u. The key idea is first to identify the version space VS of all possible linguistic summaries with truth value T greater than a user-defined threshold and then identify a personalized version space VSu ($VSu \subseteq VS$) that is specific for user u. This can be realized using two general scenarios. The first considers active learning of VSu: we first set VSu equal to VS and then generate sequentially the most informative linguistic summaries that correspond to half of the elements of VSu. User u is asked to accept or reject each most informative linguistic summary, which shrinks VSu twice. This strategy allows identifying VSu of size 1 with complexity linear with the number of qualifiers R. The second general scenario considers similarity-based learning of VSu: we first identify version spaces of similar cases to the one for which we generate linguistic summaries and then intersect these spaces to get the final VSu. If the information about the user type is unknown, it could be realized by a user providing feedback or asking for elaboration. This is the next challenge and is addressed in the next section.

Addressing this personalizing challenge results in having more catered explanations that support the user efficiently. They might feel more inclined to actually rely on and utilize an explanation when it is already more suited to them. Another aspect is the easier transition between knowledge groups, or initial evaluation for one. Moreover, due to the categorization of similar needs, it is possible to gain insights into the explanatory needs of various groups. With this, a form of metaexplanation can be formed, explaining why certain groups need certain forms of explanations. This in turn could, for example, be used by managers to get a better grasp and understanding of their team.

E. "Tell me more" explanation

As the last challenge, we consider it important to give the option to request more details to the generated explanation. While the template already covers the essence of the explanations, the user might have questions. For this, the linguistic summary needs to be adaptable to a variety of inquiries. For example, the user might not understand the strength of the quantifier and asks for clarification. A possibility here is to present them the explanation with a different quantifier so that they can see the impact the different magnitude has on the truth value. Another option is to compare drop or add another available feature-linguistic term pair in order to make the explanation more general or specific, respectively. Alternatively, the complement of the subset R of feature-linguistic term pairs can be given instead in order to give a different viewpoint of the explanation.

As a result, the user feels catered towards and should be able to fully understand the explanation. With this, they can learn from the gained insights and grow their knowledge about the explained topic.

V. DISCUSSION

In the previous section, we have outlined the challenges and their proposed solutions for developing Explainable Linguistic Summaries as an XAI tool. However, an important aspect of developing these solutions is the evaluation with respect to the end users. All of these challenges will result in aiding the user to understand the reasoning behind a decision more thoroughly and to apply the gained knowledge in their intended function. For this, in-depth user studies should be carried out. It is important to ask the end user how they perceived the explanations in terms of various aspects like understandability, usefulness, trustworthiness, and helpfulness [53], [54]. Moreover, it is also necessary to objectively test whether these explanations really have been understood and can be utilized efficiently. It has been shown that end users thought they understood the meaning of an explanation, but when put under a test, it became clear that there has been a misunderstanding of the presented information in the explanation (cf. [55]). How exactly such a task can be created might be task and/or domain dependent and also has to be investigated further. In a manufacturing company with real-time production this might be more difficult than in a less dynamic environment.

Additionally, while we believe that text explanations come naturally to us simply because of their nature, we also think that there is a potential to combine them with visuals. Therefore, we do not intend to contrast these two types of explanations, but rather have them complement each other in a meaningful way. This solution is for example also advocated by Reiter [56].

VI. CONCLUDING REMARKS

In this paper, we introduce the concept of an Explainable Linguistic Summary (XLS), meant as a linguistic summary serving as an explanation. We discuss the roadmap with the challenges that the linguistic summaries need to deal with, to become XLSs. We outline possible solutions and discuss the consequences. However, we notice that the evaluation will be the key focus in order to establish a well-defined methodology for these types of explanations. Future research will not only follow this roadmap but also investigate how these types of text explanations can be used complementary with visual explanations to help understand a certain outcome even better.

REFERENCES

- M. Cottle, W. Hoover, S. Kanwal, M. Kohn, T. Strome, and N. Treister, "Transforming health care through big data strategies for leveraging big data in the health care industry," New York, NY, Tech. Rep., 2013, http://ihealthtran. com/big-data-in-healthcare.
- [2] A. B. Arrieta, N. Díaz-Rodríguez, J. Del Ser, A. Bennetot, S. Tabik, A. Barbado, S. García, S. Gil-López, D. Molina, R. Benjamins *et al.*, "Explainable Artificial Intelligence (XAI): Concepts, taxonomies, opportunities and challenges toward responsible AI," *Information Fusion*, vol. 58, pp. 82–115, 2020.
- [3] A. Wilbik, J. M. Keller, and G. L. Alexander, "Linguistic summarization of sensor data for eldercare," in *Proceedings of the IEEE International Conference on Systems, Man, and Cybernetics (SMC 2011)*, 2011, pp. 2595–2599.
- [4] F. Portet, E. Reiter, A. Gatt, J. Hunter, S. Sripada, Y. Freer, and C. Sykes, "Automatic generation of textual summaries from neonatal intensive care data," *Artificial Intelligence*, vol. 173, no. 7-8, pp. 789–816, 2009.
- [5] OECD, The Survey of Adult Skills, 2019. [Online]. Available: https://www.oecd-ilibrary.org/content/publication/f70238c7-en
- [6] S. Chouhan, A. Wilbik, and R. Dijkman, "Explanation of anomalies in business process event logs with linguistic summaries," in *Proceedings* of WCCI 2022, 2022.
- [7] C. Wrede, M. H. M. Winands, and A. Wilbik, "Linguistic summaries as explanation mechanism for classification problems," *The 34th Benelux Conference on Artificial Intelligence and the 31th Belgian Dutch Conference on Machine Learning*, 2021.
- [8] A. Adadi and M. Berrada, "Peeking inside the black-box: A survey on explainable artificial intelligence (XAI)," *IEEE Access*, vol. 6, pp. 52 138–52 160, 2018.
- [9] D. Gunning, M. Stefik, J. Choi, T. Miller, S. Stumpf, and G.-Z. Yang, "XAI — Explainable Artificial Intelligence," *Science Robotics*, vol. 4, no. 37, p. eaay7120, 2019.
- [10] J. M. Alonso, C. Castiello, L. Magdalena, and C. Mencar, *Explainable Fuzzy Systems Paving the Way from Interpretable Fuzzy Systems to Explainable AI Systems*, ser. Studies in Computational Intelligence. Springer, 2021.
- [11] J. M. Alonso, A. Ramos-Soto, E. Reiter, and K. van Deemter, "An exploratory study on the benefits of using natural language for explaining fuzzy rule-based systems," in *Proc. of the IEEE International Conference* on *Fuzzy Systems*, 2017.
- [12] A. Mayn and K. van Deemter, "Towards generating effective explanations of logical formulas: Challenges and strategies," in Proc. of the Workshop on Interactive Natural Language Technology for Explainable Artificial Intelligence (NL4XAI) at the International Conference on Natural Language Generation (INLG), Dublin, Ireland, 2020.
- [13] I. Stepin, J. M. Alonso, A. Catala, and M. Pereira, "Generation and evaluation of factual and counterfactual explanations for decision trees and fuzzy rule-based classifiers," in *Proc. of the IEEE World Congress* on Computational Intelligence, 2020.
- [14] I. Stepin, J. M. Alonso, A. Catala, and M. Pereira-Fariña, "A Survey of Contrastive and Counterfactual Explanation Generation Methods for Explainable Artificial Intelligence," *IEEE Access*, vol. 9, pp. 11974 – 12 001, 2021.
- [15] K. Budzynska and S. Villata, "Argument mining," *The IEEE Intelligent Informatics Bulletin*, vol. 17, pp. 1–7, 2016.

- [16] A. Rago, O. Cocarascu, and F. Toni, "Argumentation-based recommendations: Fantastic explanations and how to find them," in *Proc. of the International Joint Conference on Artificial Intelligence (IJCAI)*, 2018, pp. 1949–1955.
- [17] C. Sierra, R. L. de Mántaras, and S. J. Simoff, "The argumentative mediator," in *Proc. of the European Conference on Multi-Agent Systems* (EUMAS) and the International Conference on Agreement Technologies (AT), Valencia, Spain, 2016, pp. 439–454.
- [18] M. W. Caminada, R. Kutlak, N. Oren, and W. W. Vasconcelos, "Scrutable plan enactment via argumentation and natural language generation," in *Proceedings of the 2014 International Conference on Autonomous Agents and Multi-Agent Systems*, 2014, pp. 1625–1626.
- [19] C. Hennessy, A. Bugarin, and E. Reiter, "Explaining Bayesian Networks in Natural Language: State of the Art and Challenges," in Proc. of the Workshop on Interactive Natural Language Technology for Explainable Artificial Intelligence (NL4XAI) at the International Conference on Natural Language Generation (INLG), Dublin, Ireland, 2020.
- [20] M. Pereira-Fariña and A. Bugarín, "Content determination for natural language descriptions of predictive Bayesian Networks," in *Proc. of the Conference of the European Society for Fuzzy Logic and Technology* (*EUSFLAT*). Atlantis Press, 2019, pp. 784–791.
- [21] J. Sevilla, "Explaining data using causal Bayesian Networks," in Proc. of the Workshop on Interactive Natural Language Technology for Explainable Artificial Intelligence (NL4XAI) at the International Conference on Natural Language Generation (INLG), Dublin, Ireland, 2020.
- [22] J. Forrest, S. Sripada, W. Pang, and G. Coghill, "Towards making NLG a voice for interpretable machine learning," in *Proc. of the International Conference on Natural Language Generation (INLG)*. Tilburg University, The Netherlands: Association for Computational Linguistics, 2018, pp. 177–182.
- [23] E. Mariotti, J. M. Alonso, and A. Gatt, "Towards harnessing natural language generation to explain black-box models," in *Proc. of the Workshop on Interactive Natural Language Technology for Explainable Artificial Intelligence (NL4XAI) at the International Conference on Natural Language Generation (INLG)*, Dublin, Ireland, 2020.
- [24] R. R. Selvaraju, M. Cogswell, A. Das, R. Vedantam, D. Parikh, and D. Batra, "Grad-cam: Visual explanations from deep networks via gradient-based localization," *International Journal of Computer Vision*, vol. 128, no. 2, p. 336–359, Oct 2019.
- [25] S. M. Lundberg and S.-I. Lee, "A unified approach to interpreting model predictions," in *Proceedings of the 31st International Conference on Neural Information Processing Systems*, ser. NIPS'17. Red Hook, NY, USA: Curran Associates Inc., 2017, p. 4768–4777.
- [26] M. Ribeiro, S. Singh, and C. Guestrin, ""Why should I trust you?": Explaining the predictions of any classifier," 2016, pp. 97–101.
- [27] A. Shrikumar, P. Greenside, and A. Kundaje, "Learning important features through propagating activation differences," in *Proceedings of the 34th International Conference on Machine Learning - Volume 70*, ser. ICML'17. JMLR.org, 2017, p. 3145–3153.
- [28] R. R. Yager, "A new approach to the summarization of data," *Information Sciences*, vol. 28, pp. 69–86, 1982.
- [29] J. Kacprzyk and S. Zadrozny, "Computing with words is an implementable paradigm: Fuzzy queries, linguistic data summaries, and natural-language generation," *IEEE Transactions on Fuzzy Systems*, vol. 18, no. 3, pp. 461–472, 2010.
- [30] L. A. Zadeh, "A computational approach to fuzzy quantifiers in natural languages," in *Computational Linguistics*. Elsevier, 1983, pp. 149–184.
- [31] M. Delgado, M. D. Ruiz, D. Sanchez, and M. A. Vila, "Fuzzy quantification: A state of the art," *Fuzzy Sets and Systems*, vol. 242, pp. 1 – 30, 2014.
- [32] J. Kacprzyk and S. Zadrożny, "Linguistic database summaries and their protoforms: toward natural language based knowledge discovery tools," *Information Sciences*, vol. 173, pp. 281–304, 2005.
- [33] J. Kacprzyk, R. R. Yager, and S. Zadrożny, "Fuzzy linguistic summaries of databases for an efficient business data analysis and decision support," in *Knowledge Discovery for Business Information Systems*, W. Abramowicz and J. Żurada, Eds. Boston: Kluwer, 2001, pp. 129– 152.
- [34] J. Kacprzyk, A. Wilbik, and S. Zadrożny, "Linguistic summarization of time series using a fuzzy quantifier driven aggregation," *Fuzzy Sets and Systems*, vol. 159, no. 12, pp. 1485–1499, 2008.
- [35] R. Castillo-Ortega, N. Marín, and D. Sánchez, "Linguistic query answering on data cubes with time dimension," *International Journal of Intelligent Systems*, vol. 26, no. 10, pp. 1002–1021, 2011.

- [36] P. S. Szczepaniak and J. Ochelska, "Linguistic summaries of standardized documents," in *Advances in Web Intelligence and Data Mining*, M. Last, P. S. Szczepaniak, Z. Volkovich, and A. Kandel, Eds. Springer Berlin Heidelberg, 2006, pp. 221–232.
- [37] D. Anderson, R. H. Luke, J. M. Keller, M. Skubic, M. Rantz, and M. Aud, "Linguistic summarization of video for fall detection using voxel person and fuzzy logic," *Computer Vision and Image Understanding*, vol. 1, no. 113, pp. 80–89, 2009.
- [38] D. Anderson, R. H. Luke, J. M. Keller, M. Skubic, M. J. Rantz, and M. A. Aud, "Modeling human activity from voxel person using fuzzy logic," *IEEE Transactions on Fuzzy Systems*, vol. 17, no. 1, pp. 39–49, 2008.
- [39] D. Anderson, R. H. Luke, E. Stone, and J. M. Keller, "Segmentation and linguistic summarization of voxel environments using stereo vision and genetic algorithms," in *Proceedings IEEE International Conference on Fuzzy Systems, World Congress on Computational Intelligence*, 2010, pp. 2756–2763.
- [40] M. Ros, M. Pegalajar, M. Delgado, A. Vila, D. T. Anderson, J. M. Keller, and M. Popescu, "Linguistic summarization of long-term trends for understanding change in human behavior," in *Proceedings of the IEEE International Conference on Fuzzy Systems, FUZZ-IEEE 2011*, 2011, pp. 2080–2087.
- [41] A. Wilbik, J. M. Keller, and J. C. Bezdek, "Linguistic prototypes for data from eldercare residents," *IEEE Transactions on Fuzzy Systems*, vol. 22, no. 1, pp. 110–123, 2013.
- [42] S. Zadrozny and J. Kacprzyk, "Summarizing the contents of web server logs: A fuzzy linguistic approach," in *Proceedings of FUZZ-IEEE 2007*, 2007.
- [43] A. Wilbik and R. M. Dijkman, "Linguistic summaries of process data," in 2015 IEEE International Conference on Fuzzy Systems (FUZZ-IEEE), 2015, pp. 1–7.
- [44] L. A. Zadeh, "Computation with imprecise probabilities," in Proceedings of the 12th International Conference Information Processing and Management of Uncertainty in Knowledge-based Systems, 2008.
- [45] A. Jain, T. Jiang, and J. M. Keller, "Impact of the shape of membership functions on the truth values of linguistic protoform summaries," in *Information Processing and Management of Uncertainty in Knowledge-Based Systems - 16th International Conference, IPMU 2016, Eindhoven, The Netherlands, June 20-24, 2016, Proceedings, Part I*, ser. Communications in Computer and Information Science, J. P. Carvalho, M. Lesot, U. Kaymak, S. M. Vieira, B. Bouchon-Meunier, and R. R. Yager, Eds., vol. 610. Springer, 2016, pp. 204–213.
- [46] N. Marín, G. Rivas-Gervilla, M. D. Ruiz, and D. Sánchez, "Formal concept analysis for the generation of plural referring expressions," *Inf. Sci.*, vol. 579, pp. 717–731, 2021.
- [47] R. Dijkman and A. Wilbik, "Linguistic summarization of event logs-a practical approach," *Information Systems*, vol. 67, pp. 114–125, 2017.
- [48] S. Corrente, J. R. Figueira, S. Greco, and R. Słowiński, "Multiple criteria decision support," *Handbook of Group Decision and Negotiation*, pp. 893–920, 2021.
- [49] T. M. Mitchell, Machine learning. McGraw-Hill, 1997.
- [50] S. Mohseni, N. Zarei, and E. D. Ragan, "A multidisciplinary survey and framework for design and evaluation of explainable ai systems," ACM Trans. Interact. Intell. Syst., vol. 11, no. 3–4, 2021.
- [51] H. Hagras and C. Wagner, "Towards the wide spread use of type-2 fuzzy logic systems in real world applications," *IEEE computational intelligence magazine*, vol. 7, no. 3, pp. 14–24, 2012.
- [52] A. Niewiadomski, "A type-2 fuzzy approach to linguistic summarization of data," *IEEE Trans. Fuzzy Syst.*, vol. 16, no. 1, pp. 198–212, 2008.
- [53] J. Aechtner, L. M. Cabrera, D. Katwal, P. Onghena, D. P. Valenzuela, and A. Wilbik, "Comparing user perception of explanations developed with XAI methods," in *IEEE International Conference on Fuzzy Systems*, *FUZZ-IEEE 2022, Padua, Italy, July 18-23, 2022.* IEEE, 2022, pp. 1–7.
- [54] M. Nauta, J. Trienes, S. Pathak, E. Nguyen, M. Peters, Y. Schmitt, J. Schlötterer, M. van Keulen, and C. Seifert, "From anecdotal evidence to quantitative evaluation methods: A systematic review on evaluating explainable ai," arXiv preprint arXiv:2201.08164, 2022.
- [55] C. Bove, J. Aigrain, M.-J. Lesot, C. Tijus, and M. Detyniecki, "Contextualising local explanations for non-expert users: An XAI pricing interface for insurance," in *Joint Proceedings of the ACM IUI 2021 Workshops*, 2021.
- [56] E. Reiter, "Text or graphics?" Dec 2016. [Online]. Available: https://ehudreiter.com/2016/12/26/text-or-graphics/