Explicative Human Activity Recognition using Adaptive Association Rule-Based Classification

Martin Atzmueller
Tilburg University (CSAI)
Tilburg, Netherlands
m.atzmuller@uvt.nl

Naveed Hayat
University of Kassel
Kassel, Germany
nha@cs.uni-kassel.de

Matthias Trojahn
Volkswagen AG
Wolfsburg, Germany
matthias.trojahn@volkswagen.de

Dennis Kroll
University of Kassel (ComTec)
Kassel, Germany
dennis.kroll@comtec.eecs.uni-kassel.de

***Abstract*—Computational social sensing is enabled by the Internet of Things at large scale. Using sensors, e. g., implemented in mobile and wearable devices, human behavior and activities can then be investigated, e. g., using according models and patterns. However, the obtained models are often not explicative, i. e., interpretable, transparent, and explanation-aware, which makes assessment and validation difficult for humans. This paper proposes a novel explicative classification approach featuring interpretable and explainable models. For this purpose, we embed a framework for building rule-based classifiers using class association rules. For evaluation, we apply two real-world datasets: One collected in the domain of personalized health using wearable sensors (accelerometers), the second one utilizing smartphone sensors for activity recognition. Our results indicate, that the proposed approach outperforms the baselines clearly, concerning both accuracy and complexity of the resulting predictive models.**

I. Introduction

With the Internet of Things (IoT), sensor data is becoming ubiquitous, which is important for the large-scale analysis of human behavior, and for building according (predictive) models. However, often such models are either "black-boxes" or rather complex, which makes it difficult to facilitate human interpretation, model transparency, and assessment, cf. [1]–[6].

In this paper, an adapted and substantially extended revision of [7], we propose a novel explicative machine learning approach utilizing *class association rules* (CARs): These are similar to standard association rules [8], however the right-hand-side of the rule is assigned a specific *class*, while the left-hand-side of the rule is made up of a conjunction of *attribute–value pairs* making up a pattern. For mining those, we apply subgroup discovery [9], [10] which identifies interesting subgroups described by specific patterns; these provide instant interpretation with respect to a *target concept of interest*, which translates to the *class* for CARs. We provide an instantiation of the CARMA framework for mining such rules cf. [7] for our explicative approach. We demonstrate its effectiveness using real-world human activity data in a classification approach, relating to both *general* activity recognition as well as to the medical domain and personalized health. The presented explicative approach leads to transparent, interpretable and explainable models, while still guaranteeing comparable (or better) accuracy compared to several standard baseline models.

978-1- 5386-1208-8/18/$31.00 © 2018 IEEE

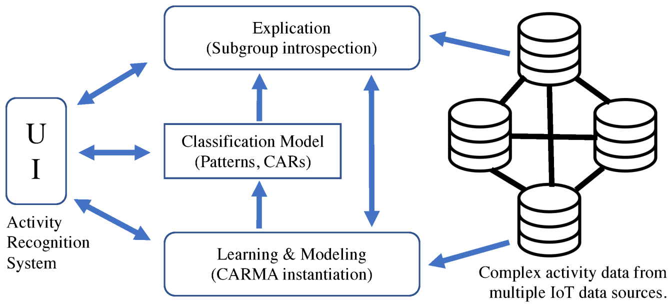

Fig. 1. **Overview: Explicative Human Activity Recognition Approach**

Figure 1 provides an overview on the proposed approach: We primarily facilitate the explication of the built model using the learned patterns and rules using subgroup introspection techniques [11], such that the CARs and patterns can be traced and inspected: That is, when interpreting the model and/or its decisions, we can utilize the rules involved in the classification for providing transparent interpretation and explanations.

Our contribution is threefold:

1) We provide a novel explicative machine learning approach, using adaptive association rule-based classification for human activity recognition [7].
2) We present an instantiation of subgroup discovery in an adaptive class association rule mining approach, targeting a classification model with low complexity but high (predictive) quality, i. e., considering metrics such as simplicity and accuracy, e. g., [12] at the same time.
3) We provide an evaluation using two real-world activity datasets. Our results indicate the efficacy of our approach by a comparison with standard symbolic classification approaches. As baselines, we utilize the (rule-based) Ripper classification algorithm as well as the (decision-tree) classifier C4.5. Our proposed method outperforms those baselines clearly.

The rest of the paper is structured as follows: Section II outlines the context of the proposed approache, discussing related work, and summarizing the applied adaptive approach for associative classification. In Section III we describe the applied datasets. Next, in Section IV we present and discuss the results of our experiments. Finally, Section V concludes the paper with a summary and outlines future work.

In: Proc. IEEE International Conference on Future IoT Technologies, IEEE Press, Boston, MA, USA, 2018

II. Method

In the following, we first discuss some related work setting the context of the proposed approach. After that, we summarize the adaptive classification approach utilizing class association rules, for which the technical implementation is described in [7] in detail.

A. Related Work

1) Classification of Sensor and Human Activity Data: Classification of activities using wearable sensors, e. g., as also integrated into mobile phones, or attached to parts of the human body, is a prominent research area, e. g., [13]–[17]. This also relates to the medical context, in particular for Parkinson's disease, e. g., Bachlin et al. [18], Mazilu et al. [19], as well as Moore et al. [20], [21]. In contrast to most of the works mentioned above, we concentrate on a set of special activities, i. e., regarding the freezing of gait for Parkinsons disease patient, as well as activities (some of which require active interaction) observed using smartphone sensors. Furthermore, we explicitly focus on explicative approaches, i. e., data mining and machine learning methods that lead to interpretable classifiers that can provide explanations for the respective models and decisions. This is a novel area, not addressed by the works mentioned above.

2) Subgroup Discovery: A broadly applicable and powerful method for descriptive and exploratory data mining is given by *subgroup discovery* [9], [10], [22], for which there are a range of efficient algorithms, e. g., [23]–[26]. In general, subgroup discovery aims at identifying patterns describing subsets of a dataset that are interesting as estimated by a quality function, e. g., [10], [22], [27]. As an example, consider a subset of a dataset that is rather specific for a certain human activity, and can be described by a logical conjunction of descriptors (e. g., attribute–value pairs). Typically, the discovered patterns are simple to interpret. This is also one of the foci of this paper. However, we go one step further than typical subgroup discovery methods that just report the interesting set of patterns to the user. For our explicative approach for human activity recognition, we will also present how to map the resulting patterns into class association rules, and how to combine these into a single classifier which is discussed in the next section.

3) Associative Classification: Associative classification integrates association rule mining into a classification approach, e. g., [28]–[30]. In the context of IoT and Wearable sensors, Hemalatha and Vijay [31] present an approach for human activity recognition and specifically fall detection using accelerometer data. For that, they mine frequent bit patterns for integration into associative classifiers. While the approach is similar regarding the pattern representation, the method proposed in this paper is more targeted using subgroup discovery for a specific class. Furthermore, it provides a refined ruleset of CARs in contrast to only focusing on frequent patterns. In addition, this procedure allows for suitable selection of a (complex) quality function for obtaining the rules, in contrast to the (simple) confidence/support-based approaches applied by association rule mining approaches.

4) Explicative Modeling: Explicative, i. e., exploratory, explainable and interpretable models have become a strong focus in the machine learning community recently, e. g., [3]–[6], [32]. Here, explanation-aware techniques [2] can be applied for a better explanation of models and their underlying mechanisms.

Explanation-aware methods are especially relevant for establishing trust in the method, e. g., [33], since the results of the method (and often also the derivation proces) can be inspected in detail in order to provide assessment and validation options. There are different flavors of explanation-aware systems. In a reconstructive explanation [34], for example, the explainer generates explanations by transforming a trace, e. g., based on intermediate results of the classification (problem solving process), into a plausible explanation story, cf. [35].

Specifically, in the area of machine learning, simple models can be targeted that feature explainability "by construction", e. g., decision trees or rule-based methods. For classification, e. g., decision paths can be followed and individual rules can then be inspected. This is the general approach taken in this paper, in contrast to methods that learn "black box" models and need to interpret them later in a second step. We apply a hybrid strategy, directly focusing on interpretable models that can then be further inspected using explanation-aware techniques in more detail, as needed.

B. CARMA*: Adaptive Associative Classification*

In the following, we summarize the applied CARMA framework for adaptive mining of class association rules and associative classification. For a detailed discussion, we refer to [7].

1) Learning: The first phase of CARMA, the learning phase constructs the predictive model: Given a set of classes, CARMA discovers class association rules for each one. In the case of a two-class problem we can optionally focus on one class only, e. g., on the minority class, since subgroup discovery is well suited for imbalanced datasets [10]. In that case, we basically focus on class association rules targeting only one particular class (e. g., the minority class). Also, for simple interpretability, we consider rules with a maximal length of the corresponding patterns. CARMA utilizes an adaptive strategy for balancing rule complexity (size) with predictive accuracy. For that purpose, a ruleset assessment function is applied complemented by a rule selection function. The ruleset assessment function can be seen as an optional heuristic. It checks, if the quality of the ruleset is good enough. If the outcome of this test is positive, CARMA continues with mining rules for the next class. Otherwise, the maximal length of a rule is increased, up to user-definable threshold. The rule selection function considers the final set of all class association rules for all classes. It selects a set of class association rules that optimize predictive power on the trainingset.

2) Classification: In the second phase, i. e., the classification phase, we utilize all the rules contained in the obtained model. Basically, given a data record, we first obtain all matching rules. Then, we aggregate their predictions using a specific rule combination strategy, e. g., cf. [36].

III. Datasets

In our experiments, we applied two datasets. For the first dataset, we utilized the publicly available *Daphnet Freezing of Gait* dataset[1] [18] in the domain of personalized health. Specifically, this dataset covers the area of detecting freezing of gait episodes for patients with Parkinson's disease. The second dataset we collected ourselves, using smartphones for detecting activity data from a variety of sensors.

A. Daphnet Freezing of Gait Dataset

The Daphnet Freezing of Gait (FoG) dataset [18] (*Daphnet/FoG*) was collected, e. g., to assess (sensor-based) methods for recognizing FoG episodes: For that, wearable sensors (accelerometers) were attached to ankle, thigh, and trunk (hip) of 10 patients with Parkinson's disease (PD). There are two classes in the dataset, *no-freezing* and *freezing (FoG)* which were assigned by manual annotation based on observational data as well as video recordings. We refer to [18] for a detailed description of the dataset and the data collection procedure.

For preprocessing the time series data, we used a sliding window with a size of four seconds for feature extraction, which is a standard parameter for activity recognition in that context, cf. [18], [19]. Furthermore, we employed data cleaning such that data records with an unknown class label were removed. This resulted in 35531 valid single activities. The final class distribution is given by 32072 data records with a label indicating "no-freezing" and 3459 data records with the label indicating "freezing (FoG)". Furthermore, we applied standard methods for feature engineering from sensor time series data in activity recognition, specifically for FoG detection and prediction cf. [19], [21]. We constructed time domain features based on several standard statistical metrics (mean, standard deviation, variance, median, range, maximum and minimum): Specifically, we considered that for each of the three accelerometer sensors (ankle, thigh and trunk), for each of their three axes (3D accelerometers), and for the squared sums of these (magnitude). This resulted in a total of 84 time domain features. Furthermore, we considered frequency domain features (based on spectral information using fast fourier transformation techniques). For that specific medical domain, we also considered special features targeting combinations of specific information from the frequency domain information, e. g., combining frequency band information from the freeze band ($3 \rightarrow 8$Hz) and the locomotor band ($0.5 \rightarrow 3$Hz), as well as the power and energy of the locomotor band, cf. [37] for details, and the freeze bands as well as the locomotor band, cf. [21] for details. As before, we computed these features for each of the three accelerometer sensors (ankle, thigh and trunk), for each of their three axes, and for the squared sums of these (magnitude). This resulted in a total of 61 frequency domain features. Therefore, for our experiments we applied 145 features in total. We applied supervised discretization [38] for deriving (nominal) selectors, i. e., attribute-value pairs, in order to utilize these for class association rule mining.

[1] http://www.ife.ee.ethz.ch/research/activity-recognition-datasets.html

TABLE I
BASELINE RESULTS USING C4.5 AND RIPPER.

Dataset	Algorithm	Accuracy	Ruleset Complexity	
			#Rules	*∅Conditions*
Daphnet/FoG	C4.5	95.16	3518	8.77
	Ripper	94.79	126	5.92
Smartphone Sensing	C4.5	69.02%	1394	6.76
	Ripper	66.87%	176	3.40

B. Smartphone Sensing Dataset

Below, we summarize the applied *smartphone sensing dataset*, and refer to cf. [7] for a detailed description. The dataset contains a diverse set of activities (classes) which included device usage as well as walking activities. We defined five scenarios that consisted of sets of different activities for data collection; these scenarios were performed by a total of 39 subjects; each scenario was repeated six times. The resulting dataset consists of a total of 3077 valid single activities, capturing data from eight different sensors.

For processing the sensor time series data, we used window-based techniques with a fixed window size of one second. This size was already proven to be efficient for walking activities [39]. We created six features per window and per sensor, cf. [7] for details on the technical implementation. Alltogether, this resulted in 116 features. As for the first dataset, we employed (supervised) discretization [38] in order to construct nominal selectors (attribute-value pairs) for subsequent class association rule mining.

IV. Evaluation

In our experiments, we utilize two baselines: Ripper [40] as a rule-based learner, and the decision-tree learner C4.5 [41] – in the JRip and J48 implementations provided by WEKA[2] [42]. For subgroup discovery we applied the BSD algorithm [25] in the implementation provided by the VIKAMINE[3] system [43]. According to the results described in [44] BSD is rather efficient for small search depths. For the evaluation, we consider (multi-class) *model accuracy* and *model complexity* as evaluation metrics. Accuracy is defined as the proportion of samples that were classified correctly. Complexity is estimated using the total number of rules contained in a rule-based model, and their average complexity. All experiments were performed in a standard 10-fold cross-validation setting.

A. Baseline Results

Table I shows the performance and complexity metrics of the baseline algorithms. For both datasets, C4.5 showed better classification performance but built more complex models with 3518 and 1394 rules, respectively. The average rule complexity scored 8.77 and 6.76, respectively. Compared to C4.5, the accuracy of Ripper is slightly lower. However, the built models are much smaller with only 126 and 176 rules, and an average rule length of 5.92 and 3.40, respectively.

[2] https://www.cs.waikato.ac.nz/ml/weka/

[3] http://www.vikamine.org

B. Experiments

In the following, we first summarize the instantiation of the experiments. After that, we present and discuss the results in detail, also concerning the explicative aspects.

1) Instantiation:

- *Ruleset assessment function*: We apply different instantiations depending on the characteristics of the applied dataset: For the *smartphone sensing* dataset which is a multi-class problem, the class distribution is not skewed; we just check, if the median of the rules' confidences is above a certain threshold τ_c. In our experiments, we applied a threshold $\tau_c = 0.5$. Regarding the *Daphnet/FoG* dataset, we focus on the minority class *freezing (Fog)* and aim to mine rather *precise* rules, i. e., those with a very high confidence. Therefore, we opt for a relatively large threshold $\tau_c = 0.99$.
- *Rule selection function*: We performed several experiments using different rule selection functions, see e. g., [28], [29]. Based on our empirical results, we chose the rule selection strategy that we proposed in [7], i. e., the CBA* selection function. It basically ensures, that for each class there is at least one rule in the final classifier. If there is none contained in the initial set of classification rules, then a default rule is added. For a detailed discussion, we refer to [7] for details.
- *Quality function*: In the multi-class case (*smartphone sensing* dataset) we selected the adjusted residuals quality function, cf. [45], which directly maps to significance criteria. For the *Daphnet/FoG* dataset, we experimented with several quality functions. Since the precision (i. e., high confidence) of the rules was the major criterion, the *relative gain* quality function, e. g., [10], showed the best results in the end.
- *Pattern Length*: We opted for interpretable patterns with a maximal length of 7 conditions.
- *Number of Patterns (TopK)*: In the evaluation, we used three different *TopK* values: 100, 200 and 500.
- *Rule combination strategy*: Based on our experiments in [7] we investigated different strategies, cf. [36], i. e., (1) *UnweightedVote*, which aggregates individual (unweighted) predictions, and chooses the class with the highest vote, (2) *LaplaceVote*, which selects according to the same principle, but the rules' weights are determined according to the Laplace value of the respective rules, (3) *BestConfidence* which classifies according to the rule with the highest confidence, and finally (4) *BestLaplace* which classifies equivently but according to the best scoring Laplace value. As described in [7], *confidence* is estimated by the respective relative frequency of the class contained in the data records covered by the respective rule, while the Laplace value $lval(r)$ of a rule r is determined by $lvar(r) = \frac{p_i^r+1}{\sum_{c_j \in C} p_j^r + |C|}$, where p_j^r (and p_i^r) are the numbers of covered examples by rule r that belong to the respective classes c_j considering all possible classes C, and class c_i of the rule, respectively.

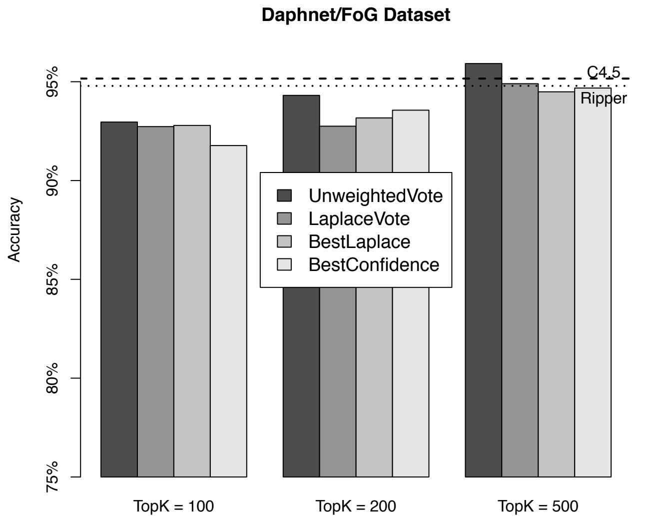

Fig. 2. Daphnet/FoG dataset: Accuracy of the proposed approach utilizing different rule combination strategies compared to the two baselines.

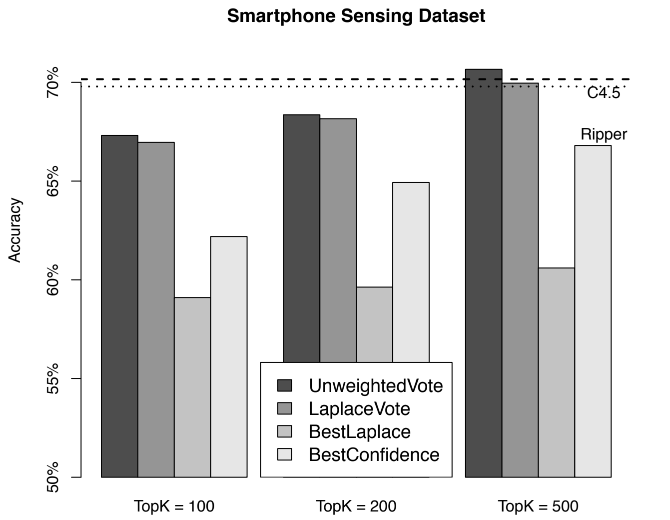

Fig. 3. Smartphone sensing dataset: Accuracy of the proposed approach utilizing different rule combination strategies compared to the two baselines.

2) Results and Discussion: For the two applied datasets, Figures 2-3 depict an overview on our experimental results, while Table II provides these in more detail.

Confirming our initial results presented in [7] for the CARMA framework, we observe that the proposed approach outperforms both baselines in accuracy as well as in complexity. In particular, the results indicate that the rule combination strategy *UnweightedVote* together with larger *TopK* values consistently achieves higher classification accuracies, which can be explained by both a more refined rule selection as well as a more precise rule combination step. The former is able to choose out of a larger set of rules in order to construct the classification model ruleset, while the latter can weight in different local patterns for estimating the final classification.

TABLE II

DETAILED EVALUATION RESULTS FOR THE DAPHNET/FOG AND THE SMARTPHONE SENSING DATASETS: THE TABLE SHOWS CLASSIFICATION ACCURACY AND COMPLEXITY (NUMBER OF RULES, AVERAGE NUMBER OF CONDITIONS) OF THE PROPOSED APPROACH DEPENDING ON DIFFERENT CHOICES OF k, AND THE APPLIED RULE COMBINATION STRATEGIES: *UnweightedVote* (UNWEIGHTED VOTING), *LaplaceVote* (VOTING USING LAPLACIAN WEIGHTS), *BestLaplace* (BEST RULE USING LAPLACE VALUE), AND *BestConfidence* (BEST RULE ACCORDING TO RULE CONFIDENCE), CF. [36].

TopK		**Daphnet/FoG**			**Smartphone Sensing**		
		Accuracy	*#Rules*	*∅ Conditions*	*Accuracy*	*#Rules*	*∅Conditions*
100	***UnweightedVote***	**92.96** %	70.8	3.96 ± 1.63	**67.31** %	345.3	2.82 ± 1.04
	LaplaceVote	92.73 %	70.2	3.97 ± 1.68	66.96 %	345.0	2.81 ± 1.04
	BestLaplace	92.79 %	71.1	3.94 ± 1.65	59.10 %	345.4	2.79 ± 0.98
	BestConfidence	91.77 %	61.9	4.00 ± 0.00	62.19 %	352.5	2.81 ± 0.96
200	***UnweightedVote***	**94.31** %	106.2	3.99 ± 0.11	**68.36** %	421.5	2.86 ± 0.97
	LaplaceVote	92.75 %	99.0	3.99 ± 0.29	68.16 %	425.0	2.90 ± 1.07
	BestLaplace	93.17 %	112.7	3.97 ± 0.37	59.63 %	423.1	2.88 ± 1.01
	BestConfidence	93.56 %	112.3	3.99 ± 0.11	64.93 %	422.8	2.89 ± 1.04
500	***UnweightedVote***	**95.92 %**	187.3	3.99 ± 0.10	**70.66 %**	517.2	3.01 ± 0.86
	LaplaceVote	94.90 %	184.1	3.99 ± 0.08	69.96 %	522.1	3.05 ± 0.96
	BestLaplace	94.49 %	187.7	3.99 ± 0.10	60.60 %	521.7	3.04 ± 0.97
	BestConfidence	94.68 %	186.3	3.99 ± 0.10	66.80 %	520.6	3.06 ± 0.97

Furthermore, we observe that the proposed approach also outperforms the baselines considering the aspect of complexity (or simplicity) of the generated models.

We summarize our experimental results as follows:

1) The proposed explicative approach utilizing the CARMA framework clearly generates less complex models than the baselines (C4.5, Ripper).
2) The presented method achieves comparable or higher accuracies with respect to the baselines (C4.5, Ripper).

Overall, these results indicate the efficacy of the proposed approach concerning accuracy and model complexity. Next, we discuss the explicative aspects of the proposed approach.

C. Explicative Classification Aspects

In order to assess the explication capabilities quantitatively, we first consider both rule complexity as well as the total number of rules. As already mentioned above, the results clearly show significant improvements compared to the baseline approaches. Rules with a low complexity are easier to understand and also provide better capabilities for statistical explanation due to less overfitting, e. g., [46], [47]. This supports, e. g., the explanation of classifications, when the set of the involved rules is presented. Therefore, this shows the efficacy concerning the interpretability and explainability. Furthermore, considering complexity measures of rule-based models, in particular the number of rules, and the average length of a rule cf. e. g., [48] it is obvious that CARMA also outperforms the baselines on that level, in particular for the C4.5 model which is the best performing baseline model.

Regarding rule coverage of the dataset, we note that larger *TopK* parameters also yield larger rule coverage fractions of the dataset, which is expected. Then, a larger number of both rules and covered examples is involved. This tackles both transparency as well as explainability since the patterns itself have more coverage and are thus backed by more examples which can be exploited in generating explanations, and inspecting patterns from different explicative perspectives, e. g., cf. [2], [11].

V. CONCLUSIONS

Human activity recognition and interpretable models for classification are prominent research directions, especially considering the ever-increasing amount of available sensor data made available by IoT contexts, e. g., in domains such as mobile assistants, health care, and Industry 4.0.

In this paper, we presented a novel approach for explicative human activity recognition using adaptive association rule-based classification. We showed the efficacy of the proposed approach in an evaluation using two real-world datasets: (i) In the context of personalized health using wearable sensors, and (ii) for activity recognition using smartphones. Our results showed, that the proposed approach outperforms the baseline approaches clearly, both in terms of accuracy and complexity of the resulting predictive models. This demonstrates the explicative capabilities of the presented approach for providing interpretable rulesets and explainable models.

For future work, we aim to investigate the applicability of the proposed approach in further domains, e. g., in the context of Industry 4.0 and according heterogeneous data. Furthermore, we plan to extend it towards complex network analytics and mining [49], and for classifying heterogeneous mobile and dynamic social media, e. g., [50], [51]. Regarding the explicative properties, we are also currently investigating the integration of prior knowledge into a hybrid explicative approach, e. g., combining the rule-based explications of the classification model with semantic knowledge for enhancing explanation-awareness. Formalizations in knowledge graphs [52], [53], for example, are an interesting direction to consider for future work.

REFERENCES

[1] M. Atzmueller, J. Baumeister, and F. Puppe, “Semi-Automatic Learning of Simple Diagnostic Scores Utilizing Complexity Measures,” *Artif Intell Med*, vol. 37, no. 1, 2006.

[2] M. Atzmueller and T. Roth-Berghofer, “The Mining and Analysis Continuum of Explaining Uncovered,” in *Proc. SGAI AI-2010*, 2010.

[3] M. T. Ribeiro, S. Singh, and C. Guestrin, “Why Should I Trust You?: Explaining the Predictions of Any Classifier,” in *Proc. SIGKDD*. ACM, 2016, pp. 1135–1144.

[4] X. Li and J. Huan, "Constructivism Learning: A Learning Paradigm for Transparent Predictive Analytics," in *Proc. SIGKDD*, 2017, pp. 285–294.
[5] O. Biran and C. Cotton, "Explanation and Justification in Machine Learning: A Survey," in *IJCAI-17 Workshop on Explainable AI*, 2017.
[6] M. Atzmueller, "Onto Explicative Data Mining: Exploratory, Interpretable and Explainable Analysis," in *Proc. Dutch-Belgian Database Day*. TU Eindhoven, Netherlands, 2017.
[7] M. Atzmueller, M. Kibanov, N. Hayat, M. Trojahn, and D. Kroll, "Adaptive Class Association Rule Mining for Human Activity Recognition," in *Proceedings of the International Workshop on Mining Ubiquitous and Social Environments (MUSE2015)*, Porto, Portugal, 2015.
[8] R. Agrawal and R. Srikant, "Fast Algorithms for Mining Association Rules," in *Proc. 20th Int. Conf. Very Large Data Bases, (VLDB)*, J. B. Bocca, M. Jarke, and C. Zaniolo, Eds. Morgan Kaufmann, 1994, pp. 487–499. [Online]. Available: citeseer.ist.psu.edu/agrawal94fast.html
[9] S. Wrobel, "An Algorithm for Multi-Relational Discovery of Subgroups," in *Proc. PKDD*, ser. LNCS, vol. 1263. Heidelberg, Germany: Springer, 1997, pp. 78–87.
[10] M. Atzmueller, "Subgroup Discovery – Advanced Review," *WIREs Data Mining and Knowledge Discovery*, vol. 5, no. 1, pp. 35–49, 2015.
[11] M. Atzmueller and F. Puppe, "A Case-Based Approach for Characterization and Analysis of Subgroup Patterns," *Journal of Applied Intelligence*, vol. 28, no. 3, pp. 210–221, 2008.
[12] M. Atzmueller, J. Baumeister, and F. Puppe, "Quality Measures and Semi-Automatic Mining of Diagnostic Rule Bases," in *Proc. INAP*, ser. LNAI, no. 3392. Heidelberg, Germany: Springer, 2005, pp. 65–78.
[13] O. D. Lara and M. A. Labrador, "A Survey on Human Activity Recognition using Wearable Sensors," *Communications Surveys & Tutorials, IEEE*, vol. 15, no. 3, pp. 1192–1209, 2013.
[14] F. Foerster, M. Smeja, and J. Fahrenberg, "Detection of Posture and Motion by Accelerometry: A Validation Study in Ambulatory Monitoring," *Computers in Human Behavior*, vol. 15, no. 5, pp. 571–583, 1999.
[15] L. Bao and S. S. Intille, "Activity Recognition from User-Annotated Acceleration Data," in *Pervasive Computing*. Springer, 2004, pp. 1–17.
[16] J. R. Kwapisz, G. M. Weiss, and S. A. Moore, "Activity Recognition using Cell Phone Accelerometers," *SIGKDD Explor. Newsl.*, vol. 12, no. 2, pp. 74–82, Dec. 2010.
[17] J. Yang, "Towards Physical Activity Diary: Motion Recognition Using Simple Acceleration Features with Mobile Phones," in *Proc. Intl. Workshop on Interactive Multimedia for Consumer Electronics*. New York, NY, USA: ACM, 2009, pp. 1–10.
[18] M. Bachlin, M. Plotnik, D. Roggen, I. Maidan, J. M. Hausdorff, N. Giladi, and G. Troster, "Wearable Assistant for Parkinson's Disease Patients with the Freezing of Gait Symptom," *IEEE Trans. on Information Technology in Biomedicine*, vol. 14, no. 2, pp. 436–446, 2010.
[19] S. Mazilu, A. Calatroni, E. Gazit, D. Roggen, J. M. Hausdorff, and G. Tröster, "Feature Learning for Detection and Prediction of Freezing of Gait in Parkinson's Disease," in *Intl. Workshop on Machine Learning and Data Mining in Pattern Recognition*. Springer, 2013, pp. 144–158.
[20] S. T. Moore, D. A. Yungher, T. R. Morris, V. Dilda, H. G. MacDougall, J. M. Shine, S. L. Naismith, and S. J. Lewis, "Autonomous Identification of Freezing of Gait in Parkinson's Disease from Lower-Body Segmental Accelerometry," *Journal of Neuroengineering and Rehabilitation*, vol. 10, no. 1, p. 19, 2013.
[21] T. T. Pham, S. T. Moore, S. J. Lewis, D. N. Nguyen, E. Dutkiewicz, A. J. Fuglevand, A. L. McEwan, and P. H. Leong, "Freezing of Gait Detection in Parkinson's Disease: A Subject-Independent Detector Using Anomaly Scores," *IEEE Transactions on Biomedical Engineering*, 2017.
[22] W. Klösgen, "Explora: A Multipattern and Multistrategy Discovery Assistant," in *Advances in Knowledge Discovery and Data Mining*. AAAI, 1996, pp. 249–271.
[23] M. Atzmueller and F. Puppe, "SD-Map - A Fast Algorithm for Exhaustive Subgroup Discovery." in *Proc. ECML/PKDD*, ser. LNAI, vol. 4213. Heidelberg, Germany: Springer, 2006, pp. 6–17.
[24] M. Atzmueller and F. Lemmerich, "Fast Subgroup Discovery for Continuous Target Concepts," in *Proc. ISMIS*, ser. LNCS, vol. 5722. Heidelberg, Germany: Springer, 2009, pp. 1–15.
[25] F. Lemmerich, M. Rohlfs, and M. Atzmueller, "Fast Discovery of Relevant Subgroup Patterns," in *Proc. FLAIRS*. Palo Alto, CA, USA: AAAI Press, 2010, pp. 428–433.
[26] F. Lemmerich, M. Becker, and M. Atzmueller, "Generic Pattern Trees for Exhaustive Exceptional Model Mining," in *Proc. ECML/PKDD*, ser. LNAI, vol. 7524. Heidelberg, Germany: Springer, 2012, pp. 277–292.
[27] L. Geng and H. J. Hamilton, "Interestingness Measures for Data Mining: A Survey," *ACM Computing Surveys*, vol. 38, no. 3, 2006.
[28] B. Liu, W. Hsu, and Y. Ma, "Integrating Classification and Association Rule Mining," in *Proc. KDD*. AAAI Press, August 1998, pp. 80–86.
[29] W. Li, J. Han, and J. Pei, "CMAR: Accurate and Efficient Classification Based on Multiple Class-Association Rules," in *Proc. ICDM*. IEEE, 2001, pp. 369–376.
[30] F. Thabtah, "A Review of Associative Classification Mining," *KER*, vol. 22, no. 1, 2007.
[31] C. S. Hemalatha and V. Vaidehi, "Associative Classification based Human Activity Recognition and Fall Detection using Accelerometer," *International Journal of Intelligent Information Technologies (IJIIT)*, vol. 9, no. 3, pp. 20–37, 2013.
[32] M. Atzmueller and E. Sternberg, "Mixed-Initiative Feature Engineering Using Knowledge Graphs," in *Proc. 9th International Conference on Knowledge Capture (K-CAP)*. ACM Press, 2017.
[33] K. Behrenbruch, M. Atzmueller, C. Evers, L. Schmidt, G. Stumme, and K. Geihs, "A Personality Based Design Approach Using Subgroup Discovery," in *Human-Centred Software Engineering*, ser. LNCS. Heidelberg, Germany: Springer, 2012, vol. 7623, pp. 259–266.
[34] M. R. Wick and W. B. Thompson, "Reconstructive Expert System Explanation," *Artif. Intell.*, vol. 54, no. 1-2, pp. 33–70, 1992.
[35] T. R. Roth-Berghofer and M. M. Richter, "On Explanation," *Künstl. Intelligenz*, vol. 22, no. 2, pp. 5–7, May 2008.
[36] J.-N. Sulzmann and J. Fürnkranz, "A Comparison of Techniques for Selecting and Combining Class Association Rules," in *Proc. LWA*, ser. Technical Report, vol. 448. Dept. of Computer Science, University of Würzburg, Germany, 2008, pp. 87–93.
[37] S. T. Moore, H. G. MacDougall, and W. G. Ondo, "Ambulatory Monitoring of Freezing of Gait in Parkinson's Disease," *Journal of neuroscience methods*, vol. 167, no. 2, pp. 340–348, 2008.
[38] U. M. Fayyad and K. B. Irani, "Multi-Interval Discretization of Continuous-Valued Attributes for Classification Learning," in *IJCAI*, 1993, pp. 1022–1029.
[39] T. Huynh and B. Schiele, "Analyzing Features for Activity Recognition," in *Proc. 2005 Joint Conf. on Smart Objects and Ambient Intelligence*. ACM, 2005, pp. 159–163.
[40] W. W. Cohen, "Fast Effective Rule Induction," in *Proc. ICML*. Morgan Kaufmann, 1995, pp. 115–123.
[41] R. Quinlan, *C4.5: Programs for Machine Learning*. San Francisco, CA, USA: Morgan Kaufmann Publishers Inc., 1993.
[42] M. Hall, E. Frank, G. Holmes, B. Pfahringer, P. Reutemann, and I. H. Witten, "The WEKA Data Mining Software: An Update," *SIGKDD Explorations*, vol. 11, no. 1, pp. 10–18, 2009.
[43] M. Atzmueller and F. Lemmerich, "VIKAMINE - Open-Source Subgroup Discovery, Pattern Mining, and Analytics," in *Proc. ECML/PKDD*, ser. LNAI, vol. 7524. Heidelberg, Germany: Springer, 2012.
[44] F. Lemmerich, M. Atzmueller, and F. Puppe, "Fast Exhaustive Subgroup Discovery with Numerical Target Concepts," *Data Mining and Knowledge Discovery*, vol. 30, pp. 711–762, 2016.
[45] A. Agresti, *An Introduction to Categorical Data Analysis*. Wiley, 2007.
[46] M. Atzmueller, J. Baumeister, and F. Puppe, "Introspective Subgroup Analysis for Interactive Knowledge Refinement," in *Proc. FLAIRS*. Palo Alto, CA, USA: AAAI Press, 2006, pp. 402–407.
[47] R. C. Holte, "Very Simple Classification Rules Perform Well on Most Commonly Used Datasets," *Machine Learning*, vol. 11, pp. 63–91, 1993.
[48] M. Atzmueller, J. Baumeister, and F. Puppe, "Semi-Automatic Learning of Simple Diagnostic Scores utilizing Complexity Measures," *Artif Intell Med*, vol. 37, no. 1, pp. 19–30, 2006.
[49] M. Atzmueller, "Data Mining on Social Interaction Networks," *Journal of Data Mining and Digital Humanities*, vol. 1, June 2014.
[50] B. Xu, A. Chin, H. Wang, L. Chang, K. Zhang, F. Yin, H. Wang, and L. Zhang, "Physical Proximity and Online User Behavior in an Indoor Mobile Social Networking Application," in *Proc. CPSCom*, 2011.
[51] M. Atzmueller, "Detecting Community Patterns Capturing Exceptional Link Trails," in *Proc. IEEE/ACM ASONAM*, 2016.
[52] M. Atzmueller, B. Kloepper, H. A. Mawla, B. Jäschke, M. Hollender, M. Graube, D. Arnu, A. Schmidt, S. Heinze, L. Schorer, A. Kroll, G. Stumme, and L. Urbas, "Big Data Analytics for Proactive Industrial Decision Support: Approaches & First Experiences in the Context of the FEE Project," *atp edition*, vol. 58, no. 9, 2016.
[53] P. Ristoski and H. Paulheim, "Semantic Web in Data Mining and Knowledge Discovery: A Comprehensive Survey," *Web Semantics*, vol. 36, pp. 1–22, 2016.