1 Introduction

The advent of the computer provided mankind with a new tool that was able to solve problems which were infeasible so far. The computer made possible the solving of difficult and extensive calculations in shorter time. Machines were increasingly improved and today they are able to make millions of calculations in the blink of an eye. Even though machines have already been calculating in speeds that outperformed human capability for a long time, they are unable to act intelligently or in a self-aware manner. So far, really intelligent and self-aware machines only exist in science fiction, such as HAL 9000 or the Terminator. But there are niches where a computer or device might learn to act "intelligently".

So far, it is not known how a computer could learn in a human way, how it could learn any arbitrary concept from its experience or its failures. Nevertheless, there are special tasks which a computer can learn to solve, e.g. trying to find treatments for a new disease from medical records or saving energy costs of a house by setting the heating apparatus according to the usage patterns of the occupants. Despite the fact that those tasks are in general quite specific and must be clearly defined, a computer can make use of its calculating speed and create new insights that the human user or even expert in that field is unaware of yet [Mit97].

The question of how a computer can be enabled to solve certain problems on its own is dealt with in *machine learning* research. The purpose of this field is to find new algorithms that let a computer create generalizing knowledge from the experience or data that was available so far.

One of the most popular and interesting problems in machine learning is *classification*. It assumes a set of examples from which each belongs to exactly one distinct class. The generalization desired from the machine learning algorithm is a mapping that is able to predict an unknown instance's true class.

Researchers from various backgrounds have developed a wealth of algorithms that are able to cope with classification problems. The perhaps most famous one might be the *neural network* that emerged from the Perceptron idea introduced by the psychologist Rosenblatt in the late 1950s [Ros58]. Others include Support Vector Machines, Logistic Regression or probabilistic approaches such as the family of Bayesian classifiers. Even though the mentioned techniques apply very different learning strategies, they all have a lack of interpretability in common. This materializes in classification decisions that can neither be understood nor explained. The reasoning mechanism within those *black box* algorithms remains untransparent. But in addition to these algorithms, symbolic approaches such as decision trees and rule-based models have also been developed: Machine learning algorithms that are considered to be interpretable. When classification decisions can be explained with the model, a human expert could intervene if the reasoning does not fit. To make those algorithms even more interpretable, researchers have conceived linguistic variants based on Zadeh's *fuzzy logic*. Fuzzy logic is an extension of the classical two-valued logic that allows intermediate degrees of truth [Zad65]. In contrast to decision trees and rule-based models the linguistic ones do not operate on the attribute values directly but on linguistically interpretable fuzzy sets instead. This meaning is attached to the fuzzy set in form of a label. Consequently, linguistic fuzzy models are very readable and understandable even without an underlying knowledge of the attribute domains. This is especially useful when a human expert is interacting with such a system.

1.1. A Brief History of Conventional and Fuzzy Rule-Based Classification

The domain of rule-based classification emerged in the 1960s, most notably with expert systems — a realm pioneered by Edward A. Feigenbaum [Fei80]. These systems were built to reproduce specific human knowledge, e.g. through inference rules. While in the beginning the rule logic had to be elicited within expert interviews, Ryszard S. Michalski developed the idea of separate-and-conquer learning, an inductive technique which is able to infer rules from data automatically. At roughly the same time, Lotfi A. Zadeh was working on a multi-valued logic which he coined fuzzy logic. Fuzzy logic was mainly intended as a formal framework of a human-like approximate reasoning and "computing with words".

From the late 1980s until the mid 1990s the realm of separate-and-conquer rule learning was studied by a large number of researchers in the machine learning community. Competition concerning efficiency and predictive quality led to significant improvements in this field. During that time a large number of heuristics and strategies to learn simple but well-classifying models emerged. A cornerstone of separate-and-conquer rule learning was developed by Johannes Fürnkranz with the IREP algorithm, which was used by William W. Cohen in the RIPPER classifier from 1995, which remains a state-of-the-art algorithm even today.

The task of classification learning was also discovered by researchers from the fuzzy set community. In 1992, Li-Xin Wang and Jerry M. Mendel published their work on how to learn a linguistic fuzzy rule-based classification model. In contrast to the discoveries from the machine learning field, Wang and Mendel used a rather simple grid-partitioning scheme that was able to yield interpretable rules but which was lacking effectivity and efficiency. Researchers from the fuzzy set community developed new algorithms as a remedy for the weakness of the initial idea. The main objective was to maintain interpretability while improving discriminative power. One of the dominating approaches since then has been the idea of using evolutionary algorithms for improving the grid-oriented linguistic fuzzy rules. In contrast to the fast and steady improvements in the machine learning community, no benchmark state-of-the-art learner within the fuzzy set community evolved. Albeit dozens of evolutionary and other hybrid algorithms for linguistic fuzzy rule learning have been published, there is none which excels. While the machine learning community designed lean but effective methods and heuristics, the fuzzy set community applied even more gargantuan approaches which were both ineffective and slow. The direct comparison of conventional with these fuzzy rule learning techniques shows clearly that the latter are not yet competitive in terms of predictive power. It is by no means an exaggeration to describe the efforts to build these gargantuan systems as a scientific dead end so far.

Basically, there has been little effort in establishing a link between the possibilities that fuzzy logic offers with the experience and efficiency from the conventional rule-learning realm. In this direction there are still opportunities for improvements for either fields.

1.2. Purpose of this Thesis

So far, the realms of conventional and fuzzy classification rule learning have been researched separately. Researchers from the machine learning community developed rule-based classification algorithms with a notable attention on the discriminative abilities and to the efficient use of processor time and memory consumption. The focus of researchers from the fuzzy logic community emphasized the aspect of interpretability when developing linguistic fuzzy rule-based classification algorithms. As a result, both communities are relatively disjunct: There are no remarkable efforts to bring the best of both worlds together. The consequence is that in the world of machine learning the fuzzy rule-based approaches are not taken seriously due to their weak predictive qualities.

This thesis will revolve around the combination of conventional rule learning methods and techniques for fuzzifying the rules. Therefore, we will introduce procedures to soften conventional rule boundaries. We will investigate the consequences of that fuzzification process and analyze the differences between the original conventional and the fuzzy version.

1.3. Contribution of this Thesis

In this thesis the following aspects will find consideration:

New data-driven fuzzy rule-based classifiers The main contribution of this thesis will be the development of novel methods for fuzzy rule-based classification. We will introduce three algorithms, two will improve on an existing state-of-the-art conventional rule learner from the separate-and-conquer realm and one will be a completely novel approach that is related to both decision tree learning and data discretization techniques.

Effective fuzzification techniques The transformation of conventional into fuzzy rules plays the key role in this work. We introduce two procedures that soften existing rule boundaries making a conventional rule fuzzy: (A) a fuzzification strategy for ordinary classification rules that do not grid-partition the numeric attributes, (B) a fuzzification strategy for rules that partition the data set into a grid.

Deep investigation into the effects of rule fuzzification This thesis encompasses serious investigations into the effects of fuzzification. In the focus are the consequences that are caused by the fuzzification process. As we will see, fuzzification comes with side effects that influence the classification decision. In a step-by-step examination the side effects will be peeled off in order to go to the very heart of fuzzification influence.

Comprehensive experimental analysis In machine learning, benefits of new algorithms are typically proven experimentally in a statistically sound way. Due to the concentration on objective, quantitive characteristics — e.g. the predictive accuracy — a comparison of two or more classifiers is rather simple. This enables researchers to compare their algorithms on a large number of data sets without making the evaluation more complicated or time consuming. Researchers are encouraged to further improve their algorithms in order to be competitive. This governance is also a kind of selection that keeps the standard high. In this work, we analyze the introduced algorithms extensively using 45 data sets from different domains.

Apart from testing the classifiers on a large testbed we will evaluate their strengths and weaknesses. To this end, we will consider classification and ranking performance for investigating the discriminative power and e.g. the number of rules and the average rule length to measure the model complexity.

The analyses will encompass thorough statistical evaluations that allow us to reach a convincing and significant conclusion.

1.4. Publications in the Context of this Thesis

Parts of this thesis have already been published in international journals or at international conferences. The following list gives an overview of the publications related to the topics of this thesis.

- J.C. Hühn and E. Hüllermeier. FR3: A fuzzy rule learner for inducing reliable classifiers. In L. Magdalena, M. Ojeda-Aciego and J.L. Verdegay, editors: *Proceedings of the 12th International Conference on Information Processing and Management of Uncertainty in Knowledge-Based Systems, IPMU*, Torremolinos (Málaga), Spain, pages 1543–1550, June 22–27, 2008.

- J.C. Hühn and E. Hüllermeier. FR3: A fuzzy rule learner for inducing reliable classifiers. *IEEE Transactions Fuzzy Systems*, 17(1):138–149, 2009.

- J.C. Hühn and E. Hüllermeier. FURIA: an algorithm for unordered fuzzy rule induction. *Data Mining and Knowledge Discovery*, 19(3)293–319, 2009.

- J.C. Hühn and S.A. Vinterbo. HELLFIRE: Learning interpretable and effective fuzzy rule-based classification models. *Fuzzy Sets and Systems*, with editor.

- J.C. Hühn and E. Hüllermeier. An analysis of the FURIA algorithm for fuzzy rule induction. In J. Koronacki, Z. Ras, S.T. Wierzchon and J. Kacprzyk, editors, *Advances in Machine Learning I: Dedicated to the memory of Professor Ryszard S. Michalski*, volume 262 of *Studies in Computational Intelligence*. Springer, Berlin, Germany, 2010.

1.5. Software Developments in the Context of this Thesis

The main part of this thesis is related to the development of fuzzy rule-based classification algorithms. All three new algorithms — FURIA, HELLFIRE and FR3 — were implemented in JAVA for the WEKA machine learning framework [WF05]. These implementations are publicly available in the software repository of the Knowledge Engineering & Bioinformatics Lab at Marburg University[1]:

`http://www.uni-marburg.de/fb12/kebi/research`

[1]The HELLFIRE implementation will be made available for download as soon as the article is accepted for publishing.

Moreover, the FURIA algorithm will be part of the official WEKA package in a contribution later than 3.7.0. Today, it is already available in the nightly build via the WEKA subversion system:

`http://www.cs.waikato.ac.nz/~ml/weka/`

1.6. Outline

This thesis will be structured following the algorithms introduced. First of all, we will provide the foundation for this dissertation in Chapter 2. This chapter will encompass the necessary theoretical and methodical background for the remaining thesis. In the next three chapters, we will introduce three different fuzzy rule-based classification algorithms: The FURIA algorithm in Chapter 3, the HELLFIRE algorithm in Chapter 4 and the FR3 algorithm in Chapter 5. In Chapter 6 we will contrast the three algorithms both from a formal but also from an experimental point of view. A comprehensive survey of related work will be given in Chapter 7. We will conclude this dissertation and give an outlook on future work in Chapter 8.