1

Introduction

Rule-based evolutionary online learning systems, often referred to as *Michigan-style learning classifier systems* (LCSs) [1], were originally inspired by the general principles of Darwinian evolution and cognitive learning. In fact, when John Holland proposed the basic LCS framework (Holland, 1976; Holland, 1977; Holland & Reitman, 1978), he actually referred to LCSs as *cognitive systems* (CSs). Inspired by stimulus-response principles in cognitive psychology, Holland designed CSs to evolve a set of production rules that convert given input into useful output. Temporary memory in the form of a *message list* was added to simulate inner mental states situating the system in the current environmental context.

Early work on LCSs confirmed the great potential of the systems for simulating animal learning and cognition as well as for real-world applications. In the first classifier system implementation, Holland and Reitman (1978) confirmed that LCSs can simulate animal behavior successfully. They evolved a representation that resulted in goal-directed, stimulus-response-based behavior satisfying multiple goals represented in resource reservoirs. Booker (1982) extended Holland's approach by experimenting with an agent that needs to avoid aversive stimuli and reach attractive stimuli. Wilson (Wilson, 1985; Wilson, 1987a) confirmed the potential of LCSs to simulate artificial animals termed *animats*—triggering the animat approach to artificial intelligence (Wilson, 1991). In brief, the approach suggests simulating animats in simulated, progressively more complex environments to understand

[1] This book is concerned with Michigan-style LCSs. These systems are online learning systems, which iteratively interact with a problem, receiving one problem instance at a time. In contrast to Michigan-style LCSs, there are Pittsburgh-style learning classifier systems (DeJong, Spears, & Gordon, 1993; Llorà & Garrell, 2001b; Llorà & Garrell, 2001a), which are batch learning systems that are much more similar to pure genetic algorithms. Despite the fundamental differences of batch-learning and the evolution of a set of solutions in these systems, a big part of the analysis in this book may be carried over—appropriately modified—to Pittsburgh-style LCSs.

learning in organisms as well as to develop highly adaptive autonomous robotic systems. Goldberg (1983) successfully applied an LCS to the control of a simulated pipeline system, confirming that LCSs are valuable learning systems for real-world applications as well.

Many of these publications were far-reaching and somewhat visionary. The LCS framework predated and inspired the now well-established reinforcement learning field (Kaelbling, Littman, & Moore, 1996; Sutton & Barto, 1998). The originally used *bucket-brigade* algorithm in LCSs (Holland, 1985) distributed reward very similar to now well-established temporal difference learning techniques, such as TD(λ) or SARSA (Sutton & Barto, 1998). The ambitious scenarios and the relation to animal learning, cognition, and robotics pointed towards research directions that remain mind boggling even today. Thus, most early LCS work laid out very interesting and challenging future research directions.

Despite these promising factors and interesting directions, the LCS framework was somewhat ahead of its time. Due to the high complexity of the systems, scalable learning of robust problem solutions could not be guaranteed. Essentially, hardly any theory was developed for an LCS system, because (1) neither learning nor convergence could be assured mathematically, (2) the learning interactions in the system appeared to be too complex and remained not well-understood, (3) the learning biases of the system were only explained intuitively, and (4) competitive applications were restricted to a somewhat limited set of smaller problems. These problematic factors led a surprisingly wide inacceptance of LCSs in the artificial intelligence and machine learning literature.

In their *"Critical review of classifier systems"*, Wilson and Goldberg (1989) pointed out several of the most important problems in the available LCSs at that time. First, it appeared that successful reward chains were hard to learn and to maintain by the means of the *bucket-brigade* algorithm in combination with the evolutionary component. Second, inappropriate bidding and payment schemes obstructed generalization, enabled overgeneralization, or prevented the formation of default hierarchies. Third, the limitations of simple classifier syntax prohibited effective processing of noisy input features, continuous problem spaces, or larger binary problem spaces. Besides these challenges, Wilson and Goldberg (1989) also mentioned the importance of developing and understanding planning and lookahead mechanisms, representations for expectations, implementations of a short-term memory, and population sizing equations.

During the subsequent LCS winter, Stewart Wilson and a few others continued to work in the LCS field. And it was Stewart Wilson who heralded an LCS renaissance with the publication of the two most influential LCS systems to date: (1) the zeroth level classifier system ZCS (Wilson, 1994) and (2) the accuracy-based classifier system XCS (Wilson, 1995).

Both classifier systems overcome many of the previously encountered challenges. The credit assignment mechanism in ZCS and XCS is directly related to the then well-understood Q-learning algorithm (Watkins, 1989) in the reinforcement learning (RL) literature so that appropriate reward estimation and propagation is ensured. Overgeneralization problems are overcome by proper fitness sharing techniques (ZCS) or the new accuracy-based fitness approach (XCS). Additionally, in XCS generalization is achieved by a niche reproduction combined with population-wide deletion, as stated in Wilson's generalization hypothesis (Wilson, 1995).

Published results suggested the competitiveness of the new LCSs (Wilson, 1994; Wilson, 1995). Solutions were found in previously unsolved maze problems that require proper generalization as well as hard Boolean function problems, such as the multiplexer problem.

Later, research focused further on the XCS system solving larger Boolean function problems (Wilson, 1998), suggesting the scalability of the system. Others focused on performance investigations in larger maze problems considering action noise and generalization (Lanzi, 1997; Lanzi, 1999a; Lanzi, 1999c).

In addition to the promising experimental results, the growth of qualitative and quantitative theoretical insights and understanding slowly gained momentum. Tim Kovacs investigated Wilson's generality hypothesis in more detail and showed that XCS strives to learn complete, accurate, and minimal representations of Boolean function problems (Kovacs, 1997). Later, Kovacs investigated the appropriate fitness approach in LCSs, contrasting a purely strength-based approach with XCS's accuracy-based approach (Kovacs, 2000; Kovacs, 2001). Finally, one of the most important questions was asked: *What makes a problem hard for XCS* (Kovacs & Kerber, 2001)? This question led to some insights on problem difficulty with respect to the optimal solution representation $[O]$. However, it remained obscured *how* XCS evolves an optimal solution as well as *which computational requirements* are necessary to successfully evolve and maintain such a solution.

Besides the new direct insights into LCSs, genetic algorithms (GAs) are now much better understood than they were back in the late 1980s. Goldberg provided a comprehensive introduction to GAs (Goldberg, 1989) and later suggested a *facetwise approach* to GA theory and design (Goldberg, 1991). The facetwise approach puts forward a modular analysis of GA components and their interactions including different selection types, structure propagation and disruption via recombination, mutation influences, or structure sustenance. The approach enabled a rigorous quantitative analysis of GA components and their interaction, which lead to a proper understanding of GA scale-up behavior and its dependence on problem structure (Goldberg, Deb, & Clark, 1992; Goldberg, Deb, & Thierens, 1993; Harik, Cantú-Paz, Goldberg, & Miller, 1997; Goldberg, 1999; Goldberg, 2002).

In addition to the *quantitative* achievements, the design decomposition also enables a rigorous *qualitative* understanding of what GAs are really searching for. Holland (1975) already hypothesized that GAs are processing *schemata*, referring to low-order attribute dependencies. GA learning success depends on the successful detection and propagation of such dependencies. However, Holland's original schema theory mainly showed the potential failure of schema processing instead of focusing on the best way to identify and propagate *useful schemata*, which are often called *building blocks* (BBs) (Holland, 1975; Goldberg, 1989; Goldberg, 2002). BBs may be characterized as lower order dependency structures (in which few attributes interact) that result in a fitness increase when set to the correct values. Attributes of a BB structure interact nonlinearly with respect to their fitness influence so that a small difference from the correct values may lead to a large difference in fitness. It should be clear that the presence of BBs highly depends on the problem at hand as well as on the chosen problem representation (Rothlauf, 2002).

It should be noted that Goldberg's facetwise analysis approach does not only facilitate system analysis and modeling but also leads to a more general system understanding and enables more effective system design (Goldberg, 2002). In the pure GA realm, for example, the GA design decomposition led to the creation of *competent GAs*—GAs that solve boundedly difficult problems quickly, accurately and reliably—including the extended compact GA (ECGA) (Harik, 1999) and the Bayesian optimization algorithm (BOA) (Pelikan, Goldberg, & Cantu-Paz, 1999) and triggering the field of estimation of distribution algorithms (EDAs) (Mühlenbein & Paaß, 1996; Pelikan, Goldberg, & Lobo, 2002; Larrañaga, 2002).

Objectives

The train of thought in this book follows a similar decomposition approach to analyze LCSs. With Wilson's powerful XCS system at hand, it establishes a rigorous understanding of XCS functioning, computational requirements, convergence properties, and generalization capabilities. The design decomposition enables us to consider evolutionary components independently so that a precise and general system analysis is possible. Meanwhile, the analysis leads us to several successfully integrated system improvements. Moreover, the proposed decomposition points towards many interesting prospective research directions including further LCS analyses and the modular and hierarchical design of more advanced LCSs.

In further detail, we first establish a rigorous understanding of LCSs, and the XCS classifier system in particular. We show which learning mechanisms can be identified, which learning biases these mechanisms cause, and how they interact. The undertaken facetwise analysis enables us to establish a fundamental theory for population sizing, problem difficulty, and learning speed. It is shown that the derived problem bounds can be used to confirm (restricted) PAC-learning capabilities of the XCS system in *k-DNF* problems. Moreover, the analysis leads us to the identification and analysis of BB-hard problems in the LCS realm. We consequently integrate competent GA recombination operators solving the BB-hard problems by making evolutionary search more effective.

Additionally, we draw connections with neural network-based function approximation techniques, combined with reinforcement learning mechanisms (Baird, 1999; Haykin, 1999; Sutton & Barto, 1998) and tabular Q-learning (Watkins, 1989). It is shown that the integration of gradient techniques improves learning reliability and accuracy in XCS. The theoretical considerations confirm that LCSs are hybrid techniques that have neural network-like interdependence properties but also tabular-like independence properties.

Besides the theoretical and mechanism-based enhancements, the book shows key-results in various problem domains including binary, nominal, and real-valued classification problems as well as multistep RL problems. Learning behavior is analyzed respecting typical problem structures and problem properties.

The lessons learned from the XCS analysis provide a broader understanding of LCSs and their interactive learning mechanisms. The analyses and experimental evaluations combined with the facetwise approach lay out a clear path towards the successful design and application of future LCS-based learning systems. The key to success is an appropriate combination of necessary learning biases comprised in the structure and type of LCS modules and their efficient interaction. Representational considerations are thereby as important as the choice of mechanisms and their effective integration.

With the gained understanding at hand, we finally propose the creation of LCS-based, cognitive learning systems that may learn interactively and incrementally a modular, distributed, and hierarchical predictive problem representation and use the representation to pursue anticipatory cognitive behavior. The proposed cognitive LCS-based structures are in accordance with Holland's original ideas but are now endowed with a modular theory on computational requirements, interactivity, learning reliability, solution accuracy and quality as the supporting backbone. The book lies out the foundations for the successful creation and application of such competent modular integrative LCS-based learning structures.

Road Map

The remainder of the book is structured as follows:

Chapter 2 provides an overview of required background knowledge. First, we introduce optimization, classification and RL problems and discuss most important structural properties, differences, and problem difficulties. Next, we provide an overview of relevant RL mechanisms. Finally, we introduce GAs focusing on Goldberg's facetwise GA decomposition approach and the aspects within most relevant for our subsequent analyses on LCSs.

Chapter 3 first gives a gentle introduction to a basic LCS system. The application to a simple toy problem illustrates the general functioning. Next, we discuss LCS theory and analysis. In particular, we propose a facetwise LCS theory approach decomposing the LCS architecture into relatively independent system facets. Each facet needs to be implemented appropriately to ensure the successful system application.

Chapter 4 introduces the system under investigation, that is, the accuracy-based classifier system XCS (Wilson, 1995). We illustrate XCS's learning behavior on exemplar toy problems, including classification and RL problems, revealing basic intuition behind XCS functioning. We then proceed to our XCS analysis.

XCS's major learning biases are investigated in Chapter 5. We show that fitness propagates accurate rules whereas generalization is achieved by a combination of subset-based reproduction and population-wide deletion. We derive a specificity equation that models the behavior of specificity in a population not influenced by fitness. Finally, we replace the previously applied proportionate selection with a subset-size dependent tournament selection mechanism ensuring reliable fitness pressure towards better classifiers.

Chapter 6 analyzes the computational requirements for solution growth and sustenance. We show that initial specificity and population size needs to be chosen adequately to ensure learning startup, minimal structural supply, relevant structural growth, and solution sustenance. With the additional learning time estimate, we can show that the computational effort scales in a low-order polynomial in problem length and solution complexity.

Next, we address solution search. Chapter 7 confirms that effective BB structure identification and processing may be necessary also in the realms of classification and RL. We introduce statistical techniques to extract evolved lower level problem structure. The gained knowledge about dependency structures is then used to mutate and recombine offspring rules more effectively, consequently solving previously hard problems successfully.

Chapter 8 applies the resulting XCS system to diverse Boolean function problems. We investigate performance in large problems, the impact of irrelevant problem features, overlapping problem subsolutions, unequally distributed subsolution complexities, and external noise. As a whole, the chapter experimentally confirms the theoretic learning bounds and supports the derived mathematical learning robustness and scalability results.

Chapter 9 applies the XCS system to real-world datamining problems as well as function approximation problems. We compare XCS's performance with several other machine learning systems in the investigated datasets. The comparison further confirms XCS's learning competence and machine learning competitiveness. We also enhance the facetwise theory to the real-valued problem domain.

Chapter 10 then investigates multistep RL problems addressing the additional challenges of reward backpropagation and distribution. The chapter shows that XCS is a competent online generalizing RL system that is able to ignore additional irrelevant problem features with additional computational effort that is linear in the number of features. The results confirm that XCS offers a robust alternative to purely neural-based RL approaches.

With pieces of the LCS puzzle then in place, Chapter 11 outlines how a similar facetwise problem approach may be helpful in the analysis of other LCSs and how LCSs can be designed using the modular theory approach. Alternative learning biases are discussed as well.

Chapter 12 then outlines how the analysis may carry over to the design of further competent and flexible LCS systems targeted to the problem at hand. In particular, we put forward the integration of LCS learning mechanisms into cognitive learning structures. Hierarchical and modular structures, anticipatory mechanisms, incremental learning, and sequential processing mechanisms are discussed. LCS mechanisms may serve as the tool-box for generating the desired structures.

Chapter 13 summarizes the major findings. The conclusions outline the next steps necessary for further LCS analysis and more competent LCS design. With the facetwise perspective on LCSs at hand, the design of Holland's originally envisioned cognitive systems may finally be within our grasp.