

Learning Classifier Systems From Principles to Modern Systems

Anthony Stein
University of Augsburg
Augsburg, Germany
anthony.stein@informatik.uni-augsburg.de

<http://gecco-2019.sigevo.org/>

Permission to make digital or hard copies of part or all of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for third-party components of this work must be honored. For all other uses, contact the Owner/Author.
GECCO '19 Companion, July 13–17, 2019, Prague, Czech Republic
© 2019 Copyright is held by the owner/author(s).
ACM ISBN 978-1-4503-6748-6/19/07.
<https://doi.org/10.1145/3319619.3323393>



Instructor

Anthony Stein is a research associate and Ph.D. candidate at the Department of Computer Science of the **University of Augsburg, Germany**. He received his B.Sc. in Business Information Systems from the University of Applied Sciences in Augsburg in 2012. He then moved to the University of Augsburg to proceed with his master's degree (M.Sc.) in computer science with a minor in information economics which he received in 2014. Within his master's thesis, he dived into the field of Learning Classifier Systems for the first time. Since then, he is a passionate follower and contributor of ongoing research in this field. His research focuses on the applicability of EML techniques in self-learning adaptive systems which are asked to act in non-stationary (i.e., real world) environments that facilitate the occurrence of knowledge gaps. Therefore, in his work he investigates the utilization of interpolation and active learning methods to change the means of how classifiers are initialized, insufficiently covered problem space niches are (proactively) filled, or adequate actions are selected. Since 2017, he is an elected organizing committee member of the International Workshop on Learning Classifier Systems (IWLCS) and since 2018 a reviewer for GECCO's EML track. He also co-organizes the Workshop Series on Autonomously Learning and Optimizing Systems (SAOS) for three years now.



What this tutorial is NOT!

- ❖ A **comprehensive** introduction to the huge field of LCS
- ❖ A **review** of all existent applications of LCS
- ❖ A in-depth **comparison** of Michigan vs. Pittsburgh LCS
- ❖ An **introduction** to the **theory** behind LCS
→ maybe in the future ;-)

What this tutorial actually is

- ❖ An **attempt** to get the audience in touch with LCS
- ❖ An **illustrative introduction** to make the LCS concept graspable
- ❖ A '**simplification**' to gain an intuition about the overarching learning framework which LCS provide
- ❖ A **starting point** to further dive into the broad field around LCS
- ❖ Therefore it is explicitly noted that...
 - we **restrict** ourselves to **Michigan-style LCS**
 - we see **abstracted views** of particular **technical details**
 - at the end **corresponding references** for a 'deeper dive' are given

Course Agenda

- ❖ Introduction
 - A Brief Definition
 - Why LCS?
 - Looking Back: LCS History
- ❖ Michigan-style Learning Classifier Systems
 - Building Blocks of LCS
 - Putting it together: A generic LCS
 - Bridging the Gap: Approaching XCS
 - Why does it learn? XCS Theory in a Nutshell
- ❖ Modern Systems
 - XCSF: Piece-wise Online Function Approximation
 - ExSTraCS: Large-scale Supervised Classification
- ❖ Summary & Conclusions
 - A Different Perspective
 - Why LCS?
 - Resources & Current Research



Course Agenda

- ❖ Introduction
 - A Brief Definition
 - Why LCS?
 - Looking Back: LCS History
- ❖ Michigan-style Learning Classifier Systems
 - Building Blocks of LCS
 - Putting it together: A generic LCS
 - Bridging the Gap: Approaching XCS
 - Why does it learn? XCS Theory in a Nutshell
- ❖ Modern Systems
 - XCSF: Piece-wise Online Function Approximation
 - ExSTraCS: Large-scale Supervised Classification
- ❖ Summary & Conclusions
 - A Different Perspective
 - Why LCS?
 - Resources & Current Research



Introduction

A Brief Definition of Learning Classifier Systems

Learning Classifier Systems (LCS) comprise a family of *flexible, evolutionary, rule-based machine learning* systems which involve a unique tandem of *local learning* and *global evolutionary optimization* of the collective models' localities.

- ❖ **Flexible**
 - Applicability: Have proven successful in a vast variety of domains
 - Extensibility: Define more a framework rather than a specific algorithm
- ❖ **Evolutionary**
 - Steady-state Niche Genetic Algorithm (GA) at their heart
 - Neo-Darwinian *Survival-of-the-Fittest* Principle: Selection, Recombination, Mutation Operators
- ❖ **Rule-based**
 - Knowledge is represented via *IF(condition)-THEN(action)* rules (aka 'classifiers')
 - *Divide-and-Conquer*: Rules partition the problem space and solve it collectively
- ❖ **Machine Learning**
 - Rules/Classifiers, i.e., their internal parameters are learnt via *stochastic gradient-based algorithms* (Widrow-Hoff delta rule, Recursive Least Squares (RLS), etc.)
 - Capable of *Reinforcement Learning (RL)*, *Supervised Learning (SL)* and *Unsupervised Learning (UL)* with only minor and straight-forward changes necessary
 - Thus, applicable to Sequential Problems, Classification, Regression, Clustering

Introduction

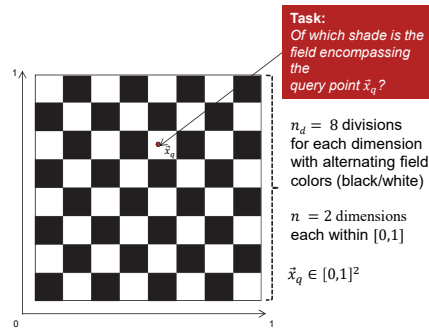
Why Learning Classifier Systems? (1/3)

- ❖ **Interpretability** by design
 - Knowledge represented by IF-THEN rules
 - Allows for explicit injection of expert knowledge
- ❖ **Complexity reduction** by design
- ❖ Online adaptivity to **dynamic learning environments**
- ❖ Inherent pressures toward **generalization**
- ❖ They are very cool ;-)
- ❖ Overarching **framework**
 - Nearly any kind of ML algorithm can be integrated
- ❖ Comparative studies confirm **competitive performance**

→ Rich body of **problem domain** and **application work** in over **40 years** of research!

Example Problem

Checkerboard Classification

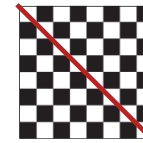


Example Problem

Checkerboard Classification

Linearly separable?

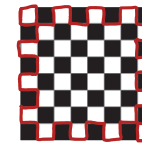
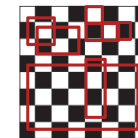
→ e.g., Linear Model, Perceptron



Non-linearly separable?

→ e.g., Multi-layer Perceptron

Problem Space Partitioning
→ LCS!



Introduction

Why Learning Classifier Systems? (2/3)

Investigated Problem Domains

- ❖ Adaptive Control (continuous and episodic)
- ❖ Uncertain Environments (Noise, Partial Observability)
- ❖ Dynamic Environments (Concept Drift/Shift)
- ❖ Data Imbalance
 - Class Imbalance
 - Sparsity regarding payoff
- ❖ High Dimensionality / Scalability
 - Exploration guidance via expert knowledge
 - Transfer Learning approaches
 - Dimensionality reduction via Autoencoders
- ❖ Complexity of underlying problem
 - Heterogeneity, Epistasis
 - Obliqueness, Curvature, Modality, etc.

Introduction

Why Learning Classifier Systems? (3/3)

Fields of Real World Application

- ❖ Gas-Pipeline Control
- ❖ Autonomous Robotics
- ❖ Robotic Kinematics
- ❖ Motion Control
- ❖ Genetics
- ❖ Biomedical Knowledge Discovery
- ❖ Medical Diagnosis
- ❖ Cognitive Modeling
- ❖ Traffic Control
- ❖ Smart Camera Networks
- ❖ Games
- ❖ ... and many more!



Introduction

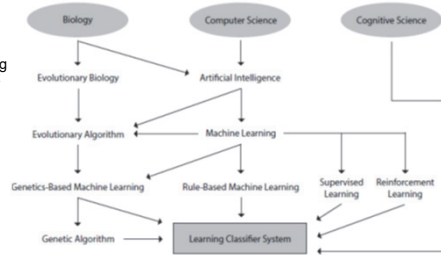
Looking Back: History of LCS*

❖ Learning Classifier System (LCS)

- ❖ In retrospect, an odd name.
- ❖ There are many machine learning systems that learn to classify but are not LCS algorithms.
- ❖ E.g. Decision trees

❖ Also referred to as...

- ❖ Rule-Based Machine Learning (RBML)
- ❖ Genetics Based Machine Learning (GBML)
- ❖ Adaptive Agents
- ❖ Cognitive Systems
- ❖ Production Systems
- ❖ Classifier System (CS, CFS)



* Image adapted from [49]

* Adapted from Urbanowicz's previous tutorials

Introduction

Looking Back: History of LCS*

1970's

- ❖ Genetic algorithms and CS-1 emerge
- ❖ Research flourishes, but application success is limited.

1980's

- ❖ LCSs are one of the earliest artificial cognitive systems – developed by **John Holland (1978)** [14].
- ❖ His work at the University of Michigan introduced and popularized the genetic algorithm.

1990's

- ❖ Holland's Vision: **Cognitive System One (CS-1)**
 - ❖ Fundamental concept of classifier rules and matching.
 - ❖ Combining a credit assignment scheme with rule discovery.
 - ❖ Function on environment with infrequent payoff/reward.

2000's

- ❖ The early work was ambitious and broad. This has led to many paths being taken to develop the concept over the following 40 years.

2010's

- ❖ CS-1 archetype would later become the basis for '**Michigan-style**' LCSs.

* Adapted from Urbanowicz's previous tutorials

Introduction

Looking Back: History of LCS*

1970's

- ❖ Pittsburgh-style algorithms introduced by **Smith** in Learning Systems One (LS-1) [35]

1980's

- ❖ LCS subtypes appear: Michigan-style vs. Pittsburgh-style
- ❖ Holland adds reinforcement learning to his system.
- ❖ Term 'Learning Classifier System' adopted.
- ❖ Research follows Holland's vision with limited success.
- ❖ Interest in LCS begins to fade.

1990's

- ❖ **Booker** suggests niche-acting GA (in [M]) [5]
- ❖ **Holland** introduces bucket brigade credit assignment [15]

2000's

- ❖ Interest in LCS begins to fade due to inherent algorithm complexity and failure of systems to behave and perform reliably

2010's

* Adapted from Urbanowicz's previous tutorials

Introduction

Looking Back: History of LCS*

1970's

- ❖ **Frey & Slate** present an LCS with predictive accuracy fitness rather than payoff-based strength [11]
- ❖ **Riolo** introduces CFCS2, setting the scene for Q-learning like methods and anticipatory LCSs [34]

1980's

- ❖ **Wilson** introduces simplified LCS architecture with his Zeroth-level Classifier System (ZCS), a strength-based system [59]

1990's

- ❖ **REVOLUTION!**
- ❖ Simplified LCS algorithm architecture with ZCS
- ❖ XCS is born: First reliable and more comprehensible LCS
- ❖ First classification and robotics applications (real-world)

2000's

- ❖ **Wilson** revolutionizes LCS algorithms with accuracy-based rule fitness in his Extended Classifier System (XCS) [60]

2010's

- ❖ **Holmes** applies LCS to problems in epidemiology [16]
- ❖ **Stolzmann** introduces Anticipatory Classifier Systems (ACS) [44]

* Adapted from Urbanowicz's previous tutorials

Introduction

Looking Back: History of LCS*

1970's

- ❖ **Wilson** introduces XCS for function approximation (XCSF) [64]
- ❖ **Kovacs** explores a number of practical and theoretical LCS questions [21,22]
- ❖ **Bernadó-Mansilla** introduce sUpervised Classifier System (UCS) for supervised learning [4]

1980's

- ❖ **Bull** explores LCS theory in simple systems [6]
- ❖ **Bacardit** introduces two Pitt-style LCS systems GAssist and BioHEL with emphasis on data mining and improved scalability to larger datasets [1,2]

1990's

- ❖ **Holmes** introduces EpiXCS for epidemiological learning. Paired with the first LCS graphical user interface to promote accessibility and ease of use [17]
- ❖ **Butz** introduces first online learning visualization for function approximation
- ❖ **Lanzi & Loiacono** explore computed actions

2000's

- ➔ LCS algorithm specializing in supervised learning and data mining start appearing
- ➔ LCS scalability becomes a central research theme
- ➔ Increasing interest in epidemiological and bioinformatics
- ➔ Facet-wise theory and applications

2010's

* Adapted from Urbanowicz's previous tutorials

Introduction

Looking Back: History of LCS*

1970's

- ❖ **Franco & Bacardit** explored GPU parallelization of LCS for scalability.
- ❖ **Urbanowicz & Moore** introduced statistical and visualization strategies for knowledge discovery in an LCS [53]. Also explored use of 'expert knowledge' to efficiently guide GA [55], introduced attribute tracking for explicitly characterizing heterogeneous patterns [54,57].

1980's

- ❖ **Browne and Iqbal** explore new concepts in reusing building blocks (i.e., code fragments) . Solved the 135-bit multiplexer reusing building blocks from simpler multiplexer problems [19].

1990's

- ❖ **Bacardit** successfully applied BioHEL to large-scale bioinformatics problems also exploring visualization strategies for knowledge discovery [3].
- ❖ **Urbanowicz** introduced ExSTraCS for supervised learning [51,56]. Applied ExSTraCS to solve the 135-bit multiplexer directly.

2000's

- ➔ Increased interest in supervised learning applications persists.
- ➔ Emphasis on solution interpretability and knowledge discovery.
- ➔ Scalability improving – 135-bit multiplexer solved!
- ➔ GPU interest for computational parallelization.
- ➔ Broadening research interest from American & European to include Australasian & Asian.

2010's

* Adapted from Urbanowicz's previous tutorials

Introduction

Looking Back: History of LCS*

1970's

1980's

1990's

2000's

2010's

~40 years of LCS research has...

- ❖ Clarified understanding.
- ❖ Produced algorithmic descriptions.
- ❖ Determined 'sweet spots' for run parameters.
- ❖ Delivered understandable 'out of the box' code.
- ❖ Demonstrated LCS algorithms to be...
 - ❖ Flexible
 - ❖ Widely applicable
 - ❖ Uniquely functional on particularly complex problems.

* Adapted from Urbanowicz's previous tutorials

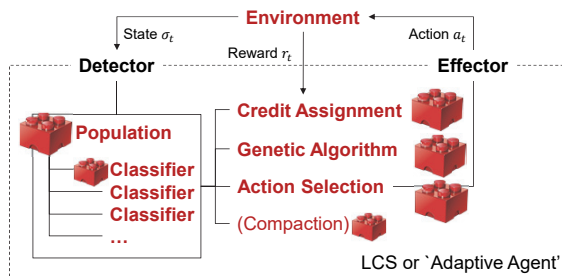
Course Agenda

- ❖ Introduction
 - ✓ A Brief Definition
 - ✓ Why LCS?
 - ✓ Looking Back: LCS History
- ❖ **Michigan-style Learning Classifier Systems**
 - Building Blocks of LCS
 - Putting it together: A generic LCS
 - Bridging the Gap: Approaching XCS
 - Why does it learn? XCS Theory in a Nutshell
- ❖ Modern Systems
 - XCSF: Piece-wise Online Function Approximation
 - ExSTraCS: Large-scale Supervised Classification
- ❖ Summary & Conclusions
 - A Different Perspective
 - Why LCS?
 - Resources & Current Research



Michigan-style LCS

Building Blocks of a Learning Classifier System

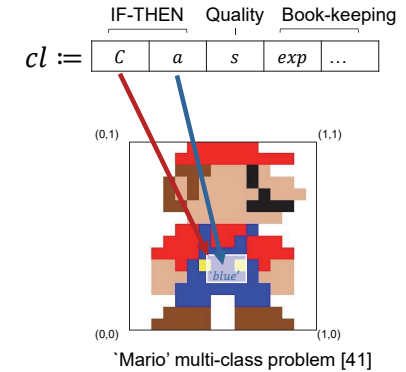


Michigan-style LCS

BBs of LCS: Classifier

Classifier cl

- ❖ IF-THEN rule
 - Condition $cl.C$
 - Action $cl.a$
- ❖ Condition $cl.C$ encodes input subspace $cl.C \subseteq X$
 - Conditions of cl 's are **not disjoint!**
- ❖ Rule strength $cl.s$, e.g.,
 - Predicted Payoff
 - Prediction Accuracy
- ❖ Book-keeping parameters
 - Experience
 - Niche size
 - Numerosity
 - etc.



* dot-notation denotes reference to parameters of specified classifier cl

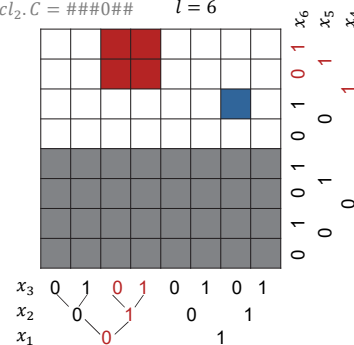
Michigan-style LCS

BBs of LCS: Classifier's Condition

Ternary Encoded Condition

- ❖ Encodes **schema** within problem's input/state space
- ❖ For **binary input spaces** \mathbb{B}^l
- ❖ One bit of input instance covered by one symbol in the condition
- ❖ Symbol from ternary alphabet $\Sigma = \{0, 1, \#\}$
 - '#' serves as don't care / wildcard
- ❖ Condition is concatenation of symbols
 - $C := (c_1, \dots, c_l), c_i \in \{0, 1, \#\}$
- ❖ Condition also encodes chromosome for the GA
- ❖ Example Problems:
 - k-Multiplexer, Majority-On, Parity, etc.

$cl_1.C = 01\#11\#$ $cl_3.C = 110101$
 $cl_2.C = \#\#\#0\#\#$ $l = 6$

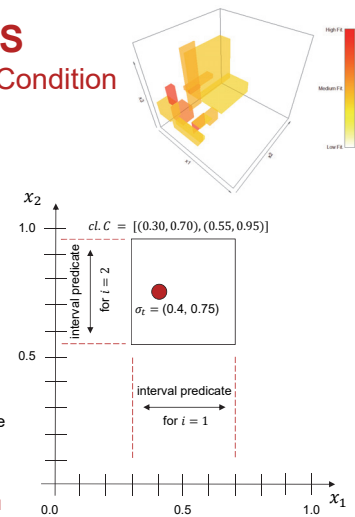


Michigan-style LCS

BBs of LCS: Classifier's Condition

Interval-based Condition

- ❖ Encodes **subspace** within problem's input/state space
- ❖ **Real-valued input spaces** \mathbb{R}^d
- ❖ One dimension $i = 1, \dots, d$ of an input instance is covered by one **interval predicate** in C
 - i -th interval predicate (l_i, u_i)
 - Lower bound l_i , upper bound u_i
 - Ordered vs. unordered Bound
- ❖ C is concatenation of intervals
 - $C := [(l_1, u_1), \dots, (l_d, u_d)], l_i, u_i \in \mathbb{R}$
- ❖ Each bound is one gene in chromosome
- ❖ Example inputs:
 - **Continuous values** e.g., Traffic flows at intersections, Sensory data
 - **Nominal** (gender, blood group) or **ordinal features** (age, salary, etc.)

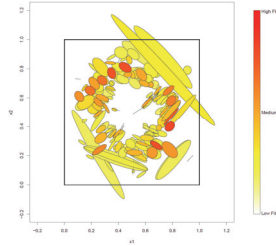
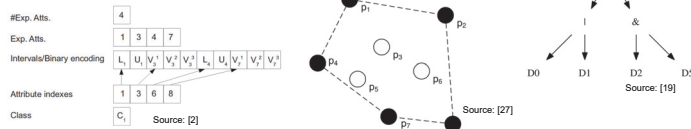


Michigan-style LCS

BBs of LCS: Classifier's Condition

Many more condition alphabets

- ❖ Hyperellipsoids (e.g., [9])
 - Covariance Matrix representation
 - Explicit geometric representation
- ❖ S-expressions / Code Fragments [19]
- ❖ Convex Hulls [27]
- ❖ Mixed Discrete-Continuous Attribute List Knowledge Representation (ALKR) [2]
- ❖ Neural Networks [7], etc.

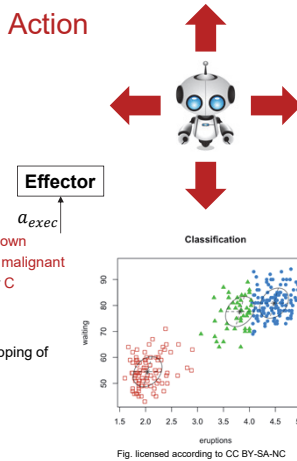


Michigan-style LCS

BBs of LCS: Classifier's Action

Discrete Actions

- ❖ Depends on the learning task
 - Reinforcement Learning: Action
 - Classification: Class/Endpoint
 - Regression: No action needed!
- ❖ Examples:
 - Robot navigation: Turn **left**, **right**, **up**, **down**
 - Medical diagnosis: Tumor is **benign** or **malignant**
 - Traffic light control: Signal **plan A**, **B** or **C**
- ❖ Large action spaces A
 - Each rule maintains a single action
 - Many rules needed for a complete mapping of the state-action-space
- ❖ Continuous Actions
 - Selection turns out difficult
 - But: Approaches do exist

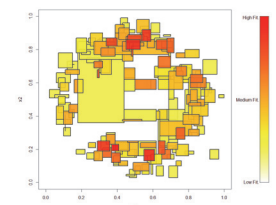
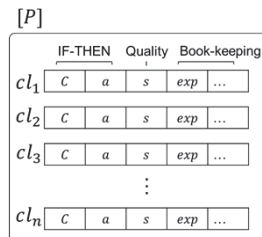


Michigan-style LCS

BBs of LCS: Population

Population $[P]$

- ❖ The **set of all rules/classifiers**
- ❖ Constitutes **knowledge base**
- ❖ Entirety of $cl \in [P]$ collectively makes up the **global model**
- ❖ Contains many **transient rules**
- ❖ Contains $n \leq N$ classifiers
 - N is a critical **hyperparameter**
 - Single classifier can **subsume** others \rightarrow **numerosity** $cl.num$
 - Size of $[P]$ is limited s.t. $\sum_{cl \in [P]} cl.num \leq N$
- ❖ $[P]$ usually starts 'tabula rasa'
- ❖ Can be initialized a priori
 - Randomly
 - Expert Knowledge / Default rules

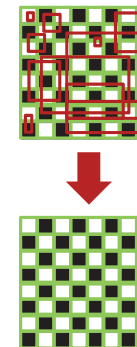


Michigan-style LCS

BBs of LCS: Compaction

Distillation of $[P]$

- ❖ Not necessary for learning success!
- ❖ Increases inference speed and comprehensibility of model
- ❖ Removes transient rules from $[P]$
 - \rightarrow Smaller collection of 'predictive' rules
- ❖ Different approaches, e.g.,
 - Condensation [60]
 - Greedy compaction [9]
 - Quick Rule Filtering [47]
- ❖ Typically applied at the end of learning or after convergence
- ❖ Up to ~90 % smaller size of $[P]$
- ❖ But only marginal increase in prediction error



Michigan-style LCS

BBs of LCS: Action Selection

Action Selection

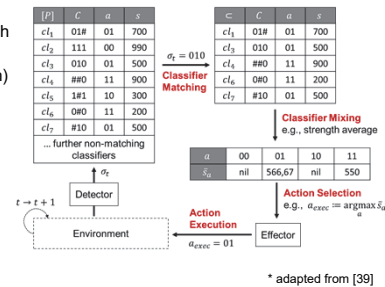
- ❖ The actual 'inference' step
- ❖ Chooses the action/prediction at each time step / for each situation
- ❖ Aka **Policy** $\pi: S \rightarrow A$ (from RL domain)
- ❖ More generally referred to as

Performance Component

- (1) Classifier **Matching** → determines niche!
- (2) Classifier **Mixing** → collective solution!
- (3) Action **Selection**
- (4) Action **Execution**

- ❖ Handles **Exploration vs. Exploitation** trade-off, e.g.,

- Interleaving random/greedy selection
- ϵ -greedy policy
- Purely explore and exploit afterwards



* adapted from [39]

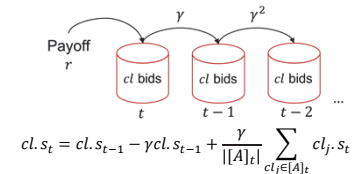
Michigan-style LCS

BBs of LCS: Credit Assignment

Credit Assignment

- ❖ Aka **Reinforcement Component**
- ❖ **Learning** comes into play
- ❖ Reward signal from environment
 - Immediate reward → may be 0
 - Delayed payoff → goal reached, 1000
- ❖ Single-step vs. Multi-step
- ❖ Correct / Incorrect Action Selection
- ❖ Reward / Punish
- ❖ Problem: Long action sequences
- ❖ Which classifiers to reinforce / attenuate?
- ❖ Early 'stage-setting' classifiers
- ❖ Adapts selected classifiers' learnable parameters, i.e., strength $cl.s$
- ❖ Updates book-keeping parameters

The early algorithm: (Implicit) **Bucket Brigade** [15,59]



The modern approach: **Temporal Difference Learning**

$$cl.s_t = cl.s_{t-1} + \beta(r_{t-1} + \gamma \max_a \bar{s}_a - cl.s_{t-1})$$

Immediate reward r_{t-1} +
current max. strength → back-up
New estimate - old estimate → TD

* Classifiers cl that were in $[A]$ of the previous cycle are updated here!

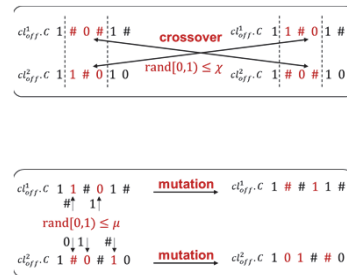
Michigan-style LCS

BBs of LCS: Genetic Algorithm

Genetic Algorithm

- ❖ Aka **Discovery Component**
- ❖ Steady-state Niche GA
- ❖ Periodic execution
- ❖ Optimizes coverage of the input space
- ❖ Usually, only conditions are altered
 - However, action mutation exists
- ❖ Fitness measure
 - Strength $cl.s$ in ZCS and older variants
 - Relative accuracy $cl.\kappa'$ in XCS and descendants (XCSF, UCS, ExSTraCS)
- ❖ Hyperparameters
 - Mutation rate μ
 - Crossover probability χ
 - Selection mechanism (Roulette-wheel vs. Tournament)
 - GA activation threshold θ_{GA}

Ternary Case



* adapted from [39]

Michigan-style LCS

BBs of LCS: Genetic Algorithm

Genetic Algorithm

- ❖ Still, steady-state niche GA
- ❖ Still, periodic execution
- ❖ Still, optimizes coverage of the input space
- ❖ Same fitness measure
- ❖ Additional hyperparameter
- ❖ Mutation spread m_0

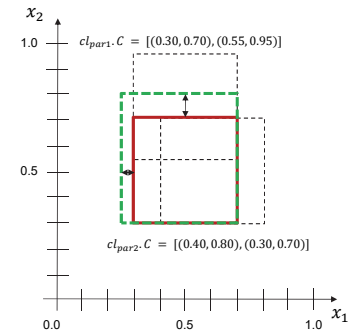
1st offspring after crossover:

$$cl_{off1}^1.C = [(0.30, 0.70), (0.30, 0.70)]$$

1st offspring after mutation:

$$cl_{off1}^2.C = [(0.25, 0.70), (0.30, 0.80)]$$

Real-valued case

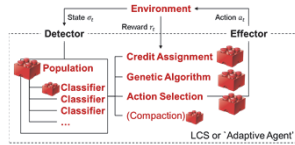


Michigan-style LCS

Putting all together



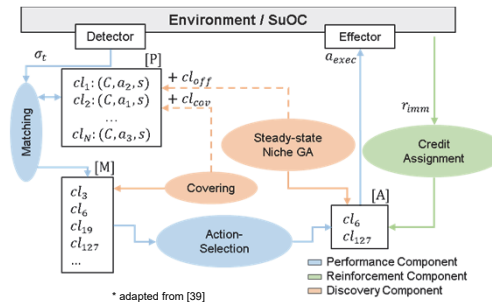
- ❖ Building blocks are the most basic components of LCS
- ❖ Each block can have more than one 'color'
- ❖ E.g., for credit assignment:
 - Bucket Brigade Algorithm
 - Profit Sharing Plan
 - Implicit Bucket Brigade
 - Q-Learning
 - Widrow-Hoff (single-step)
 - Linear Least Square
 - Recursive Least Square



- ❖ Select the most promising block for your problem and put it together
- ❖ → LCS provide a generic framework, not a single algorithm!

Michigan-style LCS

Putting all together: A Generic LCS



* adapted from [39]

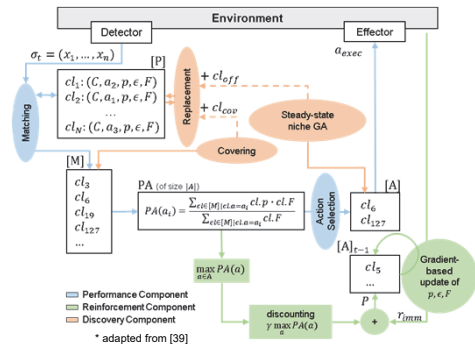
Michigan-style LCS

Bridging the Gap: Approaching XCS

- ❖ EXtended Classifier System (XCS) [60]
- ❖ Due to Stewart W. Wilson
- ❖ 'Classifier fitness based on accuracy'
- ❖ Replaces strength $cl.s$ with triplet
 - Predicted payoff $cl.p$
 - Prediction error $cl.\epsilon$
 - Fitness $cl.F$
- ❖ BBA credit assignment replaced with Q-learning-like update
- ❖ Applies niche instead of panmictic GA
 - first on $[M]$ later on $[A]$ instead of $[P]$
- ❖ Extension of the Zeroth-level Classifier System (ZCS) [59]

Michigan-style LCS

Extended Classifier System: Overview



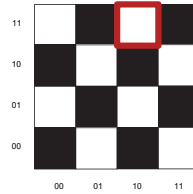
* adapted from [39]

Extended Classifier System

A quick main loop run-through

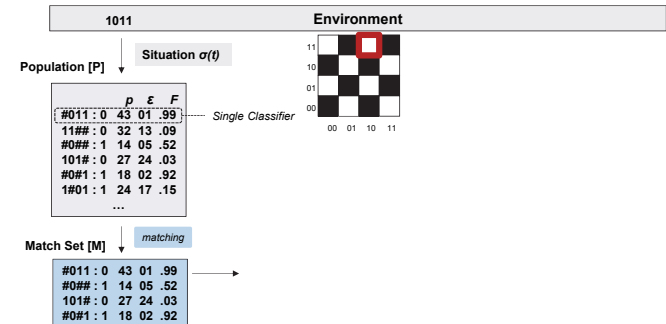
Discrete Checkerboard

- ❖ What is the situation $\sigma(t)$?
- ❖ The coordinates of the red boxed field (10,11)
- ❖ Starting horizontally: $\sigma(t)=1011$
- ❖ What are the possible actions $a \in A$?
- ❖ 'black' = 1
- ❖ 'white' = 0
- ❖ What payoff can be retrieved?
- ❖ 1000 for correct action
- ❖ 0 for wrong action



XCS Main Loop

Matching



XCS Main Loop

Matching

- ❖ At each timestep t XCS retrieves a binary string on length $n + m$
- ❖ This string is denoted as $\sigma(t) \in \{0,1\}^{n+m}$
- ❖ Example for discrete CBP ($n = 2, m = 2$ bits per dimension) and $t = 1$: $\sigma(1) = 1011$
- ❖ Each classifier maintains a **condition** C
- ❖ The conditions are encoded ternary, i.e. $C \in \{0,1,\#\}^{n+m}$
- ❖ The # symbol serves as wildcard or 'don't care' operator
- ❖ Examples of conditions: (is matching $\sigma(1)$?)
 - 1#11
 - #011
 - 01#1

Matching is the process of scanning the entire population $[P]$ for classifiers with a condition that is 'fulfilled' by the situation $\sigma(t)$

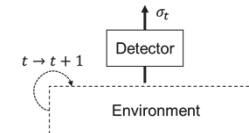
XCS Main Loop

Matching: A simple example

$[P]$	C	a	p	ϵ	F
cl_1	01#	01	700	200	0.8
cl_2	111	00	990	110	0.9
cl_3	010	01	500	500	0.5
cl_4	#00	11	900	600	0.1
cl_5	1#1	10	300	500	0.4
cl_6	0#0	11	200	50	0.9
cl_7	#10	01	500	400	0.7
... further non-matching classifiers					

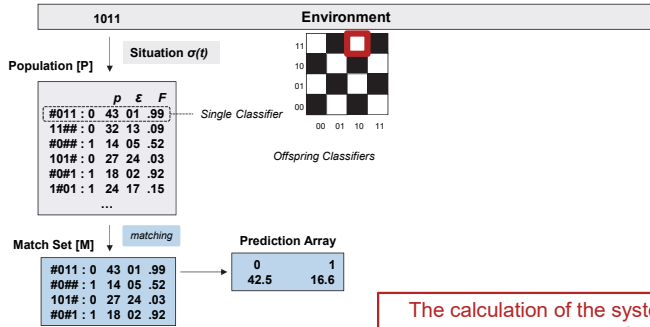
$\sigma_t = 010$
Classifier Matching

$[M]$	C	a	p	ϵ	F
cl_1	01#	01	700	200	0.8
cl_3	010	01	500	500	0.5
cl_4	#00	11	900	600	0.1
cl_6	0#0	11	200	50	0.9
cl_7	#10	01	500	400	0.7



* adapted from [39]

XCS Main Loop System Prediction



The calculation of the system prediction is the actual 'inference' step! Here, the local models are combined ('mixed') into a collective target prediction!

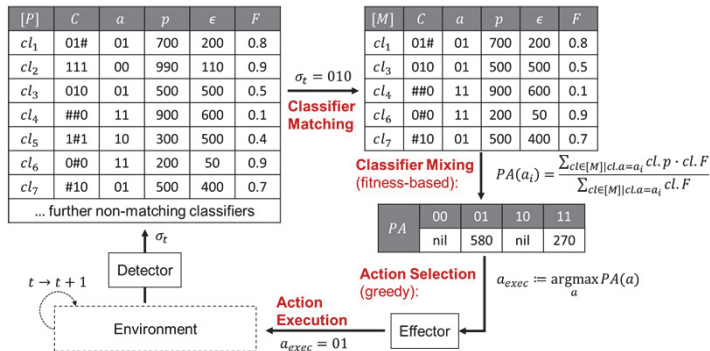
XCS Main Loop System Prediction

- ❖ The system prediction $P(a)$ is a **fitness-weighted sum of predictions** of all classifiers in $[M]$ advocating action a

$$P(a) = \frac{\sum_{cl \in [M] | cl.a=a} cl.F * cl.p}{\sum_{cl \in [M] | cl.a=a} cl.F}$$

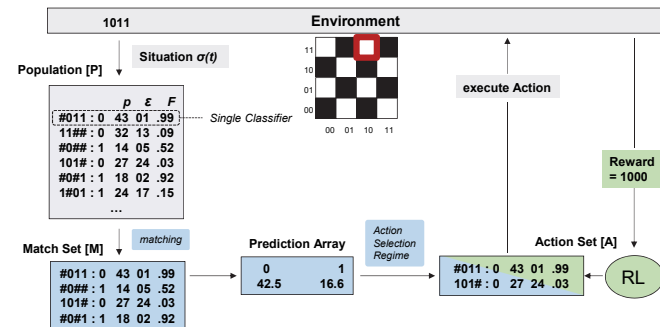
- ❖ Especially at this place, the separation of strength and accuracy becomes apparent!
- ❖ For each possible action $a \in A$ there exists one entry within the PA
- ❖ If a is not represented in $[M]$, the PA entry is *nil*

XCS Main Loop System Prediction: A simple example



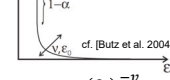
* adapted from [39]

XCS Main Loop Credit Assignment



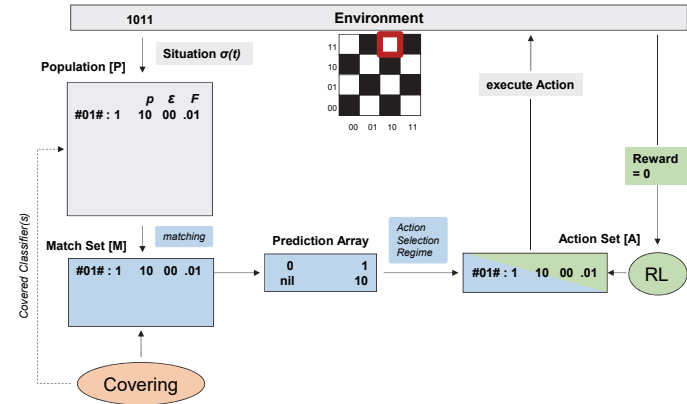
XCS Main Loop

Credit Assignment

- ❖ $\epsilon_j \leftarrow \epsilon_j + \beta(|P - p_j| - \epsilon_j)$
 - ❖ $p_j \leftarrow p_j + \beta(P - p_j)$
 - ❖ $F_j \leftarrow F_j + \beta(\kappa'_j - F_j)$, $\kappa'_j = \frac{cl_{j,\kappa} \cdot cl_{j,num}}{\sum_{cl_i \in [A]} cl_{i,\kappa} \cdot cl_{i,num}}$, $\kappa_j = \alpha \left(\frac{\epsilon_j}{\epsilon_0} \right)^{-\nu}$
- 
- ❖ β is the **learning rate** (typically set to 0.2)
 - ❖ α (often set to 0.1) and ν (usually set to 5) control **how strong accuracy decreases** when error is higher than ϵ_0
 - ❖ ϵ_0 defines the **targeted error level** of the system
 - ❖ In single-step problems, P is set to the immediate reward r_{imm}
 - ❖ Classifier parameters are updated by means of the **Widrow-Hoff (or delta) rule** in combination with the **moyenne adaptiv modifi  e (MAM)** technique

XCS Main Loop

Covering



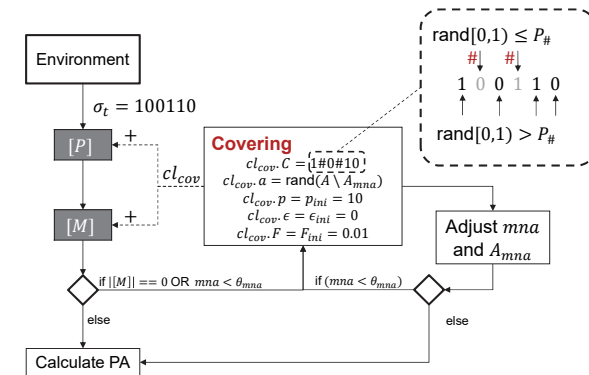
XCS Main Loop

Covering

- ❖ **Covering** is the process of **generating at least one novel classifier** that matches the current input $\sigma(t)$ whenever:
 - Match set $[M]$ is empty (i.e. no matching cl in $[P]$)
 - $[M]$ is poor, i.e. average fitness below a certain threshold
 - $[M]$ contains less than θ_{mna} distinct actions
- ❖ The condition of the covered classifier cl_{cov} is **initially set to the current input**
- ❖ Additionally, each bit is replaced by a # (for generalization purposes) with **probability $P_\#$**
- ❖ The **action** is **selected equiprobably** between actions not present in $[M]$
- ❖ Values for p, ϵ and F are set to predefined **initial values** (typically 10.0, 0.0 and 0.01, respectively)

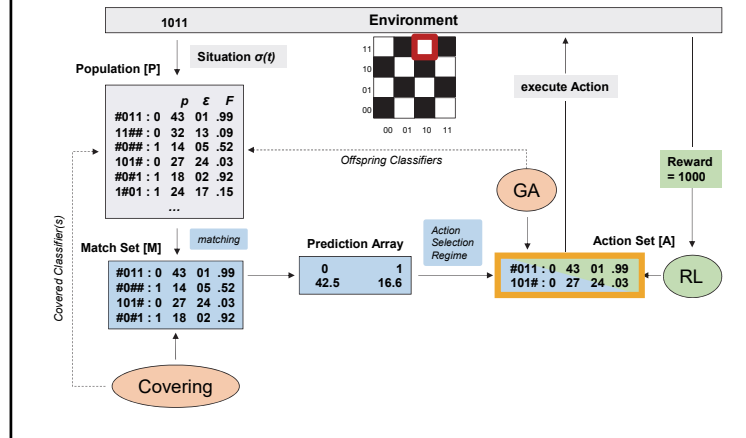
XCS Main Loop

Covering



* adapted from [39]

Genetic Algorithm

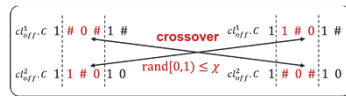


Genetic Algorithm: Invocation and Selection

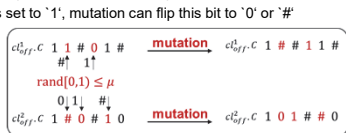
- ❖ One of the most essential parts of XCS is the incorporated steady-state niche GA (steady-state: only a small fraction of the population is replaced)
- ❖ It is triggered when the average time over all classifiers in $[A]$ since the last GA invocation is greater than θ_{GA} (often set to 12)
 - $t - \bar{\tau} > \theta_{GA}$, where $\bar{\tau} = \frac{\sum_{c \in [A]} c_{LS}}{|[A]|}$
- ❖ The GA selects two parents from $[A]$ with a probability proportional to their fitness values (roulette-wheel selection)
 - The higher a classifier's fitness, the higher the selection chance
- ❖ The selected parents are copied to generate two offspring classifiers cl_{off}^1, cl_{off}^2

Genetic Algorithm: Crossover and Mutation

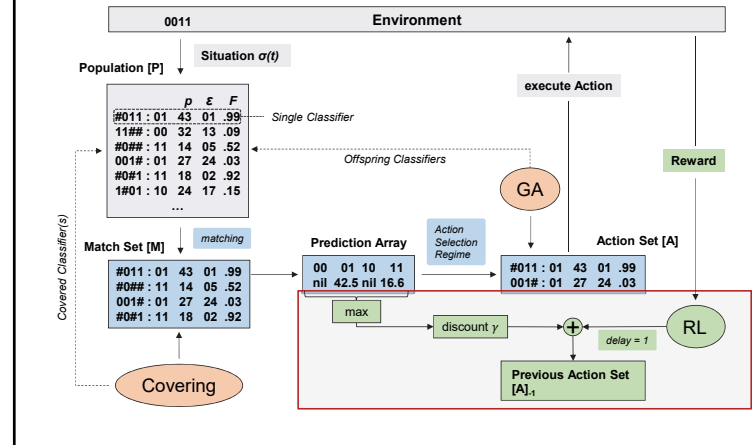
- ❖ The conditions of both cl_{off} are crossed with probability $\chi = 0.8$ (crossover operator)
 - One-point crossover: Each offspring classifier's condition is split at a certain point and switched with the other offspring classifier
 - n-point crossover: more than one point is determined for switching
 - Uniform crossover: Each value is switched with a certain probability (often 0.5)



- ❖ Afterward, each bit is flipped with probability $\mu = 0.04$ to one of the other allowed alleles (**mutation operator**)



Sequential Problem Solving (Multi-step)



XCS Main Loop

Sequential Problem Solving (Multi-step)

- ❖ r may or may not be retrieved in each step
- ❖ One has to distinguish immediate reward (r^{imm}) and total reward or payoff r at the end of a task (e.g. finally food was found)
- ❖ Update of classifier attributes is performed on the action set of the previous timestep $t - 1$ ($[A]_{-1}$)
- ❖ The maximum **system prediction** $P(a)$ from the current PA is discounted by a factor γ (usually $\gamma = 0.95$)
- ❖ Additionally, the immediate reward gained for performing the action in the previous state (of time step $t - 1$) r_{t-1}^{imm} is added (may be 0)
- ❖ This delay allows to retrieve „information from the future“
- ❖ In **single-step** environments $P = r^{imm}$
- ❖ In **multi-step** problems $P = r_{t-1}^{imm} + \gamma * \max_a PA(a)$

XCS Main Loop

Sequential Problem Solving (Multi-step)

- ❖ Single-step update of p :

$$p_j \leftarrow p_j + \beta(P - p_j)$$

- ❖ Substituting P yields us the **multi-step update formula**
- ❖ Multi-step update of p :

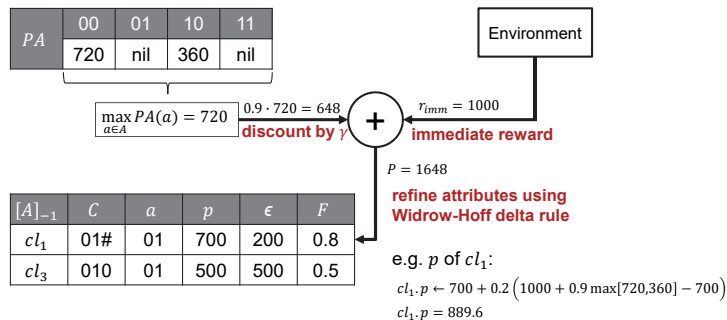
$$p_j \leftarrow p_j + \beta(r_{t-1}^{imm} + \gamma \max_a PA(a) - p_j)$$

- ❖ Do you know this update procedure from somewhere?

$$Q(s, a) \leftarrow Q(s, a) + \alpha[r + \gamma \max_{a'} Q(s', a) - Q(s, a)]$$

XCS Main Loop

Multi-step Credit Assignment: A sample calculation

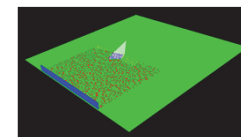
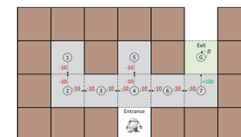


XCS Main Loop

Sequential Problem Solving (Multi-step)

Examples for multi-step environments:

- ❖ **Animat scenarios:**
 - Agent is seeking food / gold / exit / etc.
 - E.g., Woods or Maze scenarios
- ❖ Step-wise **adjustment** of a **control variable**:
 - Pan, Tilt, Zoom in Smart Camera Networks
 - Mountain Car
 - Inverse Pendulum
- ❖ **Movement decisions:**
 - 'Move to beacon' minigame in StarCraft II LE



Michigan-style LCS

XCS Theory in a Nutshell

- ❖ One disadvantage of LCS often mentioned is...

"[...] less formal understanding and a relatively small body of theoretical work [...]"

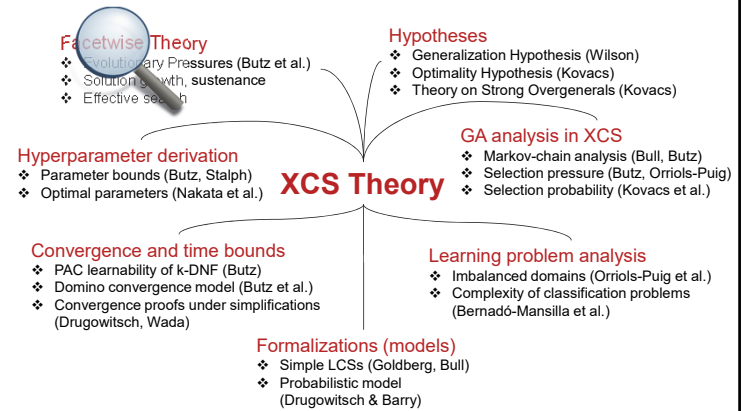
- ❖ We should put **emphasis** on "relatively"
- ❖ Sometimes experienced misconception that...

no theory  exists for LCS!

- ❖ **This is not true!**

Michigan-style LCS

XCS Theory: Much formal work already done!



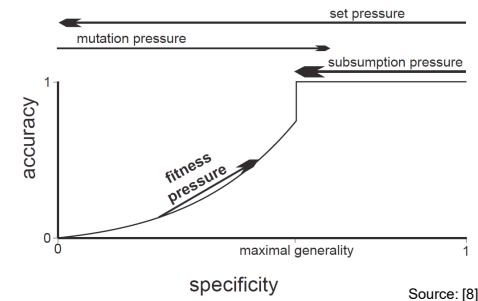
Michigan-style LCS

XCS Theory: Facetwise Approach

- ❖ **Facetwise Theory Approach** (due to Goldberg [13])
 - Proposed to analyze and understand GAs
 - Partitioning of a system into its most **relevant components**
 - Analysis in **separation**
 - Afterward, **combine** and **investigate interactions**
 - Answer questions: **What?**, **How?** and **When?**
- ❖ **Facetwise LCS Theory** (due to Butz et al. [8,10])
 - Design **evolutionary pressures** most effectively
 - Fitness guidance, parameter estimation, generalization
 - Ensure **solution growth** and **sustenance**
 - Population initialization, schema supply, growth and sustenance
 - Enable **effective solution search**
 - Mutation, recombination, local vs. global structure
 - Consider **additional challenges** in multi-step problems
 - Effective policy, problem sampling, reward propagation

Michigan-style LCS

XCS Theory: Evolutionary Pressures (or *How?*)



Michigan-style LCS

XCS Theory: Learning Bounds (or When?)

❖ Main challenges (schema and covering)

- Covering Challenge
 - Ensure coverage and GA application
 - Prevent being trapped in a **covering-deletion-cycle**
- Schema Challenge
 - Ensure that fitness pressure applies
 - From both directions: **over-general** and **over-specific** classifiers



❖ Derived bounds:

- **Covering** bound
- **Schema** bound
- **Reproductive opportunity** bound
- **Niche support** bound
- **Learning time** bound

$$\frac{-\log(1 - P(\text{cov.}))}{-\log\left(1 - \left(\frac{2 - \sigma[P]}{2}\right)^t\right)} < N$$

Covering bound, cf. [10]

❖ PAC-learnability of **k-DNF problem** confirmed for XCS with those bounds!

Course Agenda

- ❖ Introduction
 - ✓ A Brief Definition
 - ✓ Why LCS?
 - ✓ Looking Back: LCS History
- ❖ Michigan-style Learning Classifier Systems
 - ✓ Building Blocks of LCS
 - ✓ Putting it together: A generic LCS
 - ✓ Bridging the Gap: Approaching XCS
 - ✓ Why does it learn? XCS Theory in a Nutshell
- ❖ **Modern Systems**
 - XCSF: Piece-wise Online Function Approximation
 - ExSTraCS: Large-scale Supervised Classification
- ❖ Summary & Conclusions
 - A Different Perspective
 - Why LCS?
 - Resources & Current Research



Modern Systems

XCSF: Piece-wise Online Function Approximation

❖ XCS for function approximation introduced by Wilson in 2002 [64]

- Supervised learning → Actions become obsolete; only **dummy action** a_d
- Online learning → Adapt model **instance per instance**
- Local learning → Classifiers **partition the input space**; divide-and-conquer
- Evolutionary Learning → Steady-state **niche GA optimizes input space coverage**

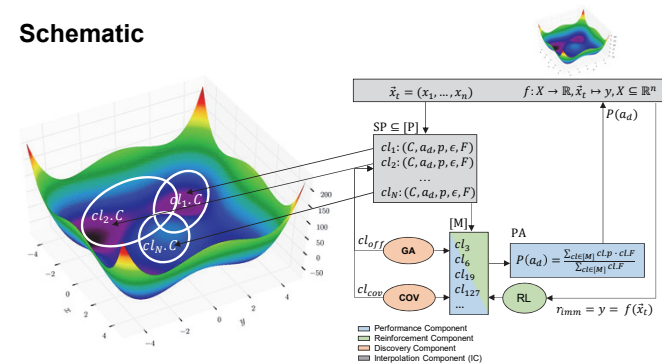
❖ **Alternative view: Evolutionary Ensemble Learner**

- XCS' algorithmic structure as a general **online ensemble learning** framework
- Classifiers as **members** of that ensemble
- No Boosting, no Bagging, more like **Stacking**
- Allows **hybrid ensemble** (cf. [26])

Modern Systems

XCSF: Piece-wise Online Function Approximation

Schematic



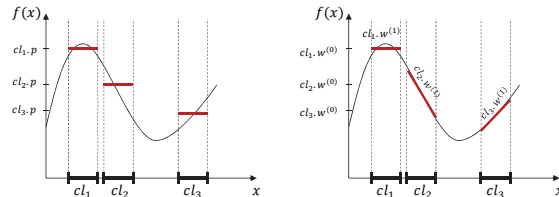
Modern Systems

XCSF: Innovations to preceding XCS(R) (1/2)

❖ Development of Classifier Prediction

- 90's: Wilson introduced ZCS and XCS as reinforcement learning algorithms
 - Classifiers cl advocate specific action $cl.a \in A$ for certain subset of states $\{\vec{x}_i\} \subseteq X$
 - Prediction attribute $cl.p$ was defined to estimate the expected reward $E[r|\vec{x}, a]$.
- 2000: XCS recognized to be well applicable to supervised learning tasks (classification).
- since 2001: Not surprisingly, it was then also used to approximate functions (regression).
- Prediction $cl.p$ was used as XCS' output
- Eventually, modeled as function $f(x) = \bar{w}^T \vec{x} + w_0$ of the current input $\vec{x} \in X$

❖ Intuition

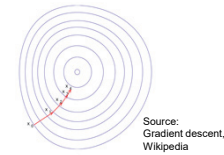


Modern Systems

XCSF: Innovations to preceding XCSR (2/2)

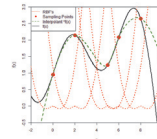
❖ Competent update procedures (cf. Lanzi et al. [24])

- Linear Least Square
- Kalman Filter
- Gain Adaptation
- Recursive Least Square



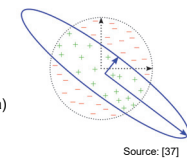
❖ Various predictors

- Polynomial approximation [25]
- Evolution Strategy [48]
- Neural Network [23]
- Support Vector Regression [29]
- RBF-Interpolation [42]



❖ Guided Mutation [37]

- Inspired by Covariance Matrix Adaptation
- Store weights for matching samples
- Assign weight < 1 for instances with high error (and vice versa)
- Guide mutation towards positively weighted instances

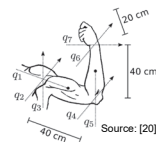


Modern Systems

XCSF: Applications

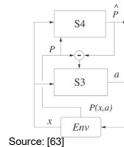
❖ Robot Kinematics

- Filtering of sensory information [20]
- Locally linear forward kinematics [38]



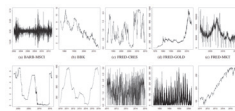
❖ Continuous action spaces [63]

- Hierarchical XCSF architectures
- e.g., Continuous Actor-Critic approach



❖ Stacking Approach for Ensemble Forecasting [36]

- Use of hybrid forecasting techniques (ARIMA, Exp. Smoothing, etc.)
- Locally learning the weights for combination of those
- Applied to different time series



Modern Systems



ExSTraCS: Large-scale Supervised Classification

- First introduced by Urbanowicz and Moore in 2014 [56]
- Conceived to tackle large-scale, complex classification problems
- Equipped with mechanisms for post-hoc Knowledge Discovery
- Proved very successful in large multiplexer problems (135-bit!)
- Focus on LCS scalability in terms of:
 - Increasing number of training instances (big data)
 - Increase in problem dimensionality (relevant features)
 - Increase in total number of features (curse of dimensionality)

❖ Open Source project (Python):

https://github.com/ryanurbs/ExSTraCS_2.0

- Visit hands-on session at IWLCS@GECCO!

ExSTraCS: Overview



* Adapted from Urbanowicz's previous tutorials

2

Pre-Processing:

ExSTraCS: Adaptive Data Management (ADM)

- Number of attributes
- Number of instances
- Location of endpoint (class)

- Automatic **shuffling** to prevent bias

- Determines **data characteristics**:

- Location of categorical attributes
 - Location of continuous attributes
 - Determines min and max ranges
 - Counts distinct values for each attribute within the training data
-
- | Attribute indexes | Class |
|-------------------|----------------|
| 1 | 1 |
| C ₁ | C ₁ |

- Automatic selection of Rule-Specificity limit (RSL)



Q

Pre-Processing: Expert Knowledge Discovery



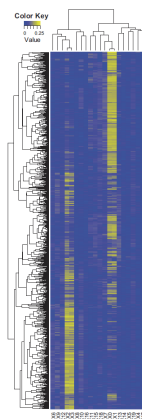
- ❖ Expert provides weights to the features/attributes
- ❖ Weights determine 'predictive value'
- ❖ Weights guide covering mechanism and GA
- ❖ Weights can be provided **manually** by **expert user**, or...
- ❖ ... **automatically** by utilizing **Relief-based attribute weighting**
 - ReliefF, SURF, SURF*, MultiSURF
 - New to ExSTraCS 2.0 → Tuned-ReliefF (TuRF)
- ❖ Introduces sort of **automated feature selection**
- ❖ But: without actual removal for knowledge discovery purposes!

* see [51,55] for more details

Attribute Tracking

ExSTraCS: Attribute Tracking und Feedback (AT&F)

- An extension to the **LCS algorithm** that allows for the **explicit characterization of heterogeneity**, and allows for the identification of **heterogeneous subject groups**.
- Akin to **long-term memory**. Experiential knowledge stored separately from the rule population that is never lost.
- Relies on learning that is both **incremental** and **supervised**.
- Stored knowledge may be fed back into LCS during learning.



* Adapted from Urbanowicz's previous tutorials

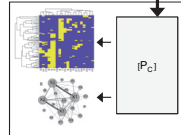
* see [54,57] for more details

Modern Systems

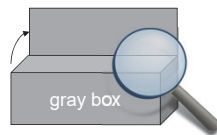
ExSTraCS: Knowledge Discovery from Output

Post-Processing: Rule Compaction

- ❖ Outputs up to 5 distinct output files
 - Final population of learned rules
 - Population metrics (train/test accuracy, etc.)
 - Attribute co-occurrence in final rules
 - Attribute tracking scores per instance
 - Summary of predictions for testing data, including votes (for further use)



- ❖ Facilitate **algorithm transparency** and **interpretability!**



* see [53,57] for more details

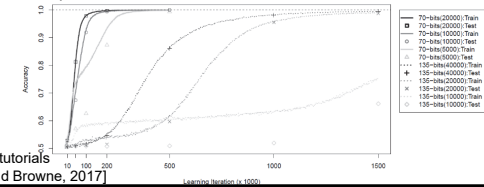
Modern Systems

ExSTraCS: Solving the 135-Multiplexer

x	Address Bits	Order of Interaction	Heterogeneous Combinations	Unique Instances	Optimal Rules [O]
6-bit	2	3	4	64	8
11-bit	3	4	8	2048	16
20-bit	4	5	16	1.05×10^6	32
37-bit	5	6	32	1.37×10^{11}	64
70-bit	6	7	64	1.18×10^{21}	128
135-bit	7	8	128	4.36×10^{40}	256

- ❖ TO SOLVE: 135-bit Multiplexer
 - All 135 features are predictive in at least some subset of the dataset.
 - Non-RBML approaches would need to include all 135 attributes together in a single model properly capturing underlying epistasis and heterogeneity.
- ❖ Few ML algorithms can make the claim that they can solve even the 6 or 11-bit multiplexer problems, let alone the 135-bit multiplexer.

ExSTraCS



* Adapted from Urbanowicz's previous tutorials

* Images adapted from [Urbanowicz and Browne, 2017]

Course Agenda

- ❖ Introduction
 - ✓ A Brief Definition
 - ✓ Why LCS?
 - ✓ Looking Back: LCS History
- ❖ Michigan-style Learning Classifier Systems
 - ✓ Building Blocks of LCS
 - ✓ Putting it together: A generic LCS
 - ✓ Bridging the Gap: Approaching XCS
 - ✓ Why does it learn? XCS Theory in a Nutshell
- ❖ Modern Systems
 - ✓ XCSF: Piece-wise Online Function Approximation
 - ✓ ExSTraCS: Large-scale Supervised Classification
- ❖ **Summary & Conclusions**
 - A Different Perspective
 - Why LCS?
 - Resources & Current Research



Summary & Conclusions

A Different (ML-centric) Perspective on LCS

- ❖ Reconsider M-style LCS as an **online ensemble learner**
- ❖ Rules = ensemble members
- ❖ Each rule constitutes a **local model / hypothesis**
- ❖ Rules are **experts** of different problem niches → **mixture of experts**
- ❖ 'Goodness' of each 'expert' determined instance-by-instance without necessity to remember → **one-pass (online) learning**
- ❖ **Modularity** (recall building block intuition) allows for **stacking**
 - Different models for local prediction (ANN, RBF, polynomials) and fitness-weighted combination = **stacked generalization**
- ❖ Learning **Classifier System** not only about classification (alone)
 - XCSF: Function Approximation = **Regression**
 - XCS(R): Sequential Decision Making = **Reinforcement Learning**
 - XCSC: Clustering = **Unsupervised Learning**
- ❖ XCSF → similarities to **Locally Weighted Projection Regression**
- ❖ XCS(R) → **generalizing Q-learner**

Summary & Conclusions

So, again: Why LCS? (ex post)

- ❖ **Flexibility** (RL, SL) and **modularity** (building blocks)
- ❖ **Interpretability** by design (condition-action rules)
- ❖ Follow **divide and conquer** principle (mixture of experts)
- ❖ Capture **complex associations** (epistasis, heterogeneity)
- ❖ Evolution as central component allows **adaptation to change** (concept drift)
- ❖ **Overarching framework** for general ML techniques
 - LCS and Deep Learning do not mutually exclude!
 - E.g., put DNNs to locally model a policy
- ❖ And (again) finally...
 - they are simply cool ;-)

Summary & Conclusions

Recent Research Directions (excerpt)

- ❖ Visual and statistical **knowledge discovery** from LCS rule sets (Urbanowicz et al. [57])
- ❖ Theoretical **hyperparameter derivation** (Nakata et al. [30,31])
- ❖ **Hierarchical** LCS and **multi-domain** learning (Liu, Browne, Xue [28])
- ❖ **Interpolation**-assisted LCS (Stein et al. [40][42][43])
- ❖ LCS with **active learning** (Stein et al. [41])
- ❖ **Algebraic formalization** of LCS (Pätzel and Hähner [32])
- ❖ ...
- ❖ → nearly all of them regularly attend GECCO!

Acknowledgements

Thanks to Ryan J. Urbanowicz for the permission to reuse parts of his previous tutorials on LCS.

Resources

- ❖ Additional Information:
 - ❖ Keep up to date with the latest LCS research
 - ❖ Get in contact with an LCS researcher
 - ❖ Contribute to the LCS community research and discussions.
- ❖ GBML Central - <http://gbml.org/>
- ❖ LCS Researcher Webpages:
 - ❖ Urbanowicz, Ryan - <http://www.ryanurbanowicz.com/>
 - ❖ Browne, Will - <http://ecs.victoria.ac.nz/Main/WillBrowne>
 - ❖ Lanzi, Pier Luca - <http://www.pierlucalanzi.net/>
 - ❖ Wilson, Stewart - <https://www.eskimo.com/~wilson/>
 - ❖ Bacardit, Jaume - <http://homepages.cs.ncl.ac.uk/jaume.bacardit/>
 - ❖ Holmes, John - <https://www.med.upenn.edu/apps/faculty/index.php/g5455356/p19936>
 - ❖ Kovacs, Tim - <http://www.cs.bris.ac.uk/home/kovacs/>
 - ❖ Bull, Larry - <http://www.cems.uwe.ac.uk/~lbull/>
- ❖ International Workshop Learning Classifier Systems (IWLCS)
- held annually at GECCO
- ❖ Mailing List: Yahoo Group: [ics-and-gbml\[at\]yahoogroups.com](mailto:ics-and-gbml[at]yahoogroups.com)

* Adapted from Urbanowicz's previous tutorials

Resources: Available Software

- ❖ Educational LCS (eLCS) – in Python.
 - <https://github.com/ryanurbs/eLCS>
 - Simple Michigan-style LCS for learning how they work and how they are implemented.
 - Code intended to be paired with first LCS introductory textbook by Urbanowicz/Browne.
- ❖ ExSTraCS 2.0 – Extended Supervised Learning LCS – in Python
 - https://github.com/ryanurbs/ExSTraCS_2.0
 - For prediction, classification, data mining, knowledge discovery in complex, noisy, epistatic, or heterogeneous problems.
- ❖ BioHEL – Bioinformatics-oriented Hierarchical Evolutionary Learning – in C++
 - <http://ico2s.org/software/biohel.html>
 - GAssist also available through this link.
- ❖ XCSLib (XCS and XCSF) (by Lanzi in C++)
 - <http://xcslib.sourceforge.net/>
- ❖ XCSF with function approximation visualization – in Java
 - [Martin Butz Chair website](#)

* Adapted from Urbanowicz's previous tutorials

Resources: LCS Review Papers & Books

Selected Review Papers:

- ❖ Pätzelt, David, Stein, Anthony, and Hähner, Jörg. "A Survey on Formal Theoretical Advances Regarding XCS." *IWLCS* (2019), under review, to appear.
- ❖ Bull, Larry. "A brief history of learning classifier systems: from CS-1 to XCS and its variants." *Evolutionary Intelligence* (2015): 1-16.
- ❖ Bacardit, Jaume, and Xavier Llorà. "Large-scale data mining using genetics-based machine learning." *Wiley Interdisciplinary Reviews: Data Mining and Knowledge Discovery* 3.1 (2013): 37-61.
- ❖ Urbanowicz, Ryan J., and Jason H. Moore. "Learning classifier systems: a complete introduction, review, and roadmap." *Journal of Artificial Evolution and Applications* 2009 (2009): 1.
- ❖ Sigaud, Olivier, and Stewart W. Wilson. "Learning classifier systems: a survey." *Soft Computing* 11.11 (2007): 1065-1078.
- ❖ Holland, John H., et al. "What is a learning classifier system?." *Learning Classifier Systems*. Springer Berlin Heidelberg, 2000. 3-32.
- ❖ Lanzi, Pier Luca, and Rick L. Riolo. "A roadmap to the last decade of learning classifier system research (from 1989 to 1999)." *Learning Classifier Systems*. Springer Berlin Heidelberg, 2000. 33-61.

Books:

- ❖ Drugowitsch, J., (2008) *Design and Analysis of Learning Classifier Systems: A Probabilistic Approach*. Springer-Verlag.
- ❖ Bull, L., Bernado-Mansilla, E., Holmes, J. (Eds.) (2008) *Learning Classifier Systems in Data Mining*. Springer
- ❖ Butz, M (2006) *Rule-based evolutionary online learning systems: A principled approach to LCS analysis and design*. Studies in Fuzziness and Soft Computing Series, Springer.
- ❖ Bull, L., Kovacs, T. (Eds.) (2005) *Foundations of learning classifier systems*. Springer.
- ❖ Kovacs, T. (2004) *Strength or accuracy: Credit assignment in learning classifier systems*. Springer.
- ❖ Butz, M. (2002) *Anticipatory learning classifier systems*. Kluwer Academic Publishers.
- ❖ Lanzi, P.L., Stolzmann, W., Wilson, S., (Eds.) (2000). *Learning classifier systems: From foundations to applications* (LNAI 1813). Springer.
- ❖ Holland, J. H. (1975). *Adaptation in natural and artificial systems*. University of Michigan Press.

* Adapted from Urbanowicz's previous tutorials

Resources: Most recent

- ❖ Textbook: 'Introduction to Learning Classifier Systems' Springer, 2017 (Urbanowicz & Brown, 2017)
- ❖ LCS Introductory Chapter: 'Reaction Learning', Chapter 7.1 in book: 'Organic Computing – Technical Systems for Survival in the Real World', Birkhäuser, 2017 (Stein, 2017)
- ❖ YouTube video on LCS:
 - ❖ Learning Classifier Systems in a Nutshell
 - ❖ Animated, narrated explanation of basic LCS concepts.
 - ❖ https://www.youtube.com/watch?v=CRqg_cZ2cJc
- ❖ LCS and Rule-Based Machine Learning Wikipedia Pages – recently updated and revised. (https://en.wikipedia.org/wiki/Learning_classifier_system)



* Adapted from Urbanowicz's previous tutorials

References

Figures

- ❖ Figure sources: All figures that have not been created by the author or indicated otherwise are free to use and taken from pixabay.com licensed according to the [Pixabay License](#)

References (1/5)

- (1) Bacardit, Jaume, et al. "Speeding-up Pittsburgh learning classifier systems: Modeling time and accuracy." *Parallel Problem Solving from Nature-PPSN VIII*. Springer Berlin Heidelberg, 2004.
- (2) Bacardit, Jaume, and Natalio Krasnogor. "A mixed discrete-continuous attribute list representation for large scale classification domains." *Proceedings of the 11th Annual conference on Genetic and evolutionary computation*. ACM, 2009.
- (3) Bacardit, Jaume, and Xavier Llorà. "Large-scale data mining using genetics-based machine learning." *Wiley Interdisciplinary Reviews: Data Mining and Knowledge Discovery* 3.1 (2013): 37-61.
- (4) Bernadó-Mansilla, Ester, and Josep M. Garrell-Guiu. "Accuracy-based learning classifier systems: models, analysis and applications to classification tasks." *Evolutionary Computation* 11.3 (2003): 209-238.
- (5) Booker, Lashon Bernard. "Intelligent behavior as an adaptation to the task environment, University of Michigan." *Ann Arbor, MI* (1982).
- (6) Bull, Larry. "A simple accuracy-based learning classifier system." *Learning Classifier Systems Group Technical Report UWELCSG03-005, University of the West of England, Bristol, UK* (2003).
- (7) Bull, Larry, and O'Hara, Toby. "Accuracy-based neuro and neuro-fuzzy classifier systems." *GECCO 2002, 905-911, Morgan Kaufmann, 2002*.
- (8) Butz, M.; Kovacs, T.; Lanzi, P. & Wilson, S., "Toward a Theory of Generalization and Learning in XCS", IEEE Transactions on Evolutionary Computation, 8, 28-46, 2004
- (9) Butz, M.; Lanzi, P. & Wilson, S. "Function Approximation With XCS: Hyperellipsoidal Conditions, Recursive Least Squares, and Compaction", IEEE Transactions on Evolutionary Computation, 12, 355-376, 2008
- (10) Butz, M. V., "Rule-based Evolutionary Online Learning Systems: A Principled Approach to LCS Analysis and Design", Springer, 2005
- (11) Frey, Peter W., and David J. Slate. "Letter recognition using Holland-style adaptive classifiers." *Machine Learning* 6.2 (1991): 161-182.
- (12) Goldberg, David E. "E. 1989. Genetic Algorithms in Search, Optimization, and Machine Learning." *Reading: Addison-Wesley* (1990).

References (2/5)

- (13) Goldberg, D. E., "Genetic Algorithms as a Computational Theory of Conceptual Design", Rzevski, G. & Adey, R. A. (Eds.), Applications of Artificial Intelligence in Engineering VI, Springer Netherlands, 1991, 3-16
- (14) Holland, J., and J. Reitman. "Cognitive systems based on adaptive agents.", *Pattern-directed inference systems* (1978).
- (15) Holland, J. "Properties of the Bucket brigade." *In Proceedings of the 1st International Conference on Genetic Algorithms*, 1-7 (1985)
- (16) Holmes, John H. "A genetics-based machine learning approach to knowledge discovery in clinical data." *Proceedings of the AMIA Annual Fall Symposium*. American Medical Informatics Association, 1996.
- (17) Holmes, John H., and Jennifer A. Sager. "The EpiXCS workbench: a tool for experimentation and visualization." *Learning Classifier Systems*. Springer Berlin Heidelberg, 2007. 333-344.
- (18) Iqbal, Muhammad, Will N. Browne, and Mengjie Zhang. "Extending learning classifier system with cyclic graphs for scalability on complex, large-scale boolean problems." *Proceedings of the 15th annual conference on Genetic and evolutionary computation*. ACM, 2013.
- (19) Iqbal, M.; Browne, W. N. & Zhang, M., "Reusing Building Blocks of Extracted Knowledge to Solve Complex, Large-Scale Boolean Problems", IEEE Transactions on Evolutionary Computation, 2014, 18, 465-480
- (20) Kneissler, J.; Stalph, P. O.; Drugowitsch, J. & Butz, M. V., "Filtering Sensory Information with XCSF: Improving Learning Robustness and Robot Arm Control Performance", Evolutionary Computation, 2014, 22, 139-158
- (21) Kovacs, Tim. "A comparison of strength and accuracy-based fitness in learning classifier systems." *School of Computer Science, University of Birmingham, Birmingham, UK* (2002).
- (22) Kovacs, Tim. "What should a classifier system learn and how should we measure it?" *Soft Computing* 6.3-4 (2002): 171-182.
- (23) Lanzi, P. L. & Loiacono, D., "XCSF with Neural Prediction", IEEE CEC, 2006, 2270-2276
- (24) Lanzi, P. L.; Loiacono, D.; Wilson, S. W. & Goldberg, D. E., "Generalization in the XCSF Classifier System: Analysis, Improvement, and Extension", *Evol. Comput.*, MIT Press, 2007, 15, 133-168
- (25) Lanzi, P. L.; Loiacono, D.; Wilson, S. W. & Goldberg, D. E., "Extending XCSF Beyond Linear Approximation", *GECCO 2005*, ACM, 2005, 1827-1834

References (3/5)

- (26) Lanzi, P. L.; Loiacono, D. & Zanini, M., "Evolving classifier ensembles with voting predictors", IEEE CEC 2008, June 1-6, 2008, Hong Kong, China, 2008, 3760-3767
- (27) Lanzi, P. L. & Wilson, S. W., "Using Convex Hulls to Represent Classifier Conditions", *GECCO 2006*, ACM, 2006, 1481-1488
- (28) Liu, Y.; Xue, B. & Browne, W. N., "Visualisation and Optimisation of Learning Classifier Systems for Multiple Domain Learning", *Simulated Evolution and Learning*, Springer International Publishing, 2017, 448-461
- (29) Loiacono, D.; Marelli, A. & Lanzi, P. L., "Support vector regression for classifier prediction", *GECCO 2007*, 2007, 1806-1813
- (30) Nakata, M.; Browne, W. N. & Hamagami, T., "Theoretical adaptation of multiple rule-generation in XCS", *GECCO 2018*, Kyoto, Japan, July 15-19, 2018, 482-489
- (31) Nakata, M.; Browne, W. N.; Hamagami, T. & Takadama, K., "Theoretical XCS parameter settings of learning accurate classifiers", *GECCO 2017*, Berlin, Germany, July 15-19, 2017, 473-480
- (32) Pätzelt, D. & Hähner, J., "An Algebraic Description of XCS", *GECCO 2018 Companion*, ACM, 2018, 1434-1441
- (33) Pätzelt, D.; Stein, A. & Hähner, J., "A Survey on Formal Theoretical Advances Regarding XCS", *GECCO 2019 Companion*, ACM, 2019, under review
- (34) Riolo, Rick L. "Lookahead planning and latent learning in a classifier system." *Proceedings of the first international conference on simulation of adaptive behavior on From animals to animats*. MIT Press, 1991.
- (35) Smith, Stephen Frederick. "A learning system based on genetic adaptive algorithms.", Dissertation, University of Pittsburgh, 1980.
- (36) Sommer, M.; Stein, A. & Hähner, J., "Local ensemble weighting in the context of time series forecasting using XCSF", *IEEE SSCI*, 2016
- (37) Stalph, P. O. & Butz, M. V., "Guided Evolution in XCSF", *GECCO 2012*, ACM, 2012, 911-918
- (38) Stalph, P. O. & Butz, M. V., "Learning local linear Jacobians for flexible and adaptive robot arm control", *GPDM*, 2012, 13, 137-157
- (39) Stein, A., "Reaction Learning", In book: *Organic Computing – Technical Systems for Survival in the Real World*, Müller-Schloer, C. & Tomforde, S. (Eds.), Birkhäuser, 2017, 287-328

References (4/5)

- (40) Stein, A.; Eymüller, C.; Rauh, D.; Tomforde, S. & Hähner, J., "Interpolation-based Classifier Generation in XCSF", *IEEE CEC*, 2016, 3990-3998
- (41) Stein, A.; Maier, R. & Hähner, J., "Toward Curious Learning Classifier Systems: Combining XCS with Active Learning Concepts", *GECCO 2017 Companion*, ACM, 2017, 1349-1356
- (42) Stein, A.; Menssen, S. & Hähner, J., "What About Interpolation? A Radial Basis Function Approach to Classifier Prediction Modeling in XCSF", *GECCO 2018*, ACM, 2018
- (43) Stein, A.; Rauh, D.; Tomforde, S. & Hähner, J., "Interpolation in the eXtended Classifier System: An Architectural Perspective", *Journal of Systems Architecture*, 75, 79-94, 2017
- (44) Stolzmann, Wolfgang. "An introduction to anticipatory classifier systems." *Learning Classifier Systems*. Springer Berlin Heidelberg, 2000. 175-194.
- (45) Stone, Christopher, and Larry Bull. "For real! XCS with continuous-valued inputs." *Evolutionary Computation* 11.3 (2003): 299-336.
- (46) Tamee, K.; Bull, L. & Pinngern, O., "Towards Clustering with XCS", *GECCO 2007*, ACM, 2007, 1854-1860
- (47) Tan, J.; Moore, J. & Urbanowicz, R., "Rapid Rule Compaction Strategies for Global Knowledge Discovery in a Supervised Learning Classifier System", *The 2018 Conference on Artificial Life: A Hybrid of the European Conference on Artificial Life (ECAL) and the International Conference on the Synthesis and Simulation of Living Systems (ALIFE)*, 2013, 110-117
- (48) Tran, T. H.; Sanza, C. & Duthen, Y., "Evolving prediction weights using evolution strategy", *GECCO 2008*, 2009-2016, 2008
- (49) Urbanowicz, Ryan J., and Jason H. Moore. "Learning classifier systems: a complete introduction, review, and roadmap." *Journal of Artificial Evolution and Applications* 2009 (2009): 1.
- (50) Urbanowicz, Ryan J., and Will Browne. "An Introduction to Learning Classifier Systems". Springer, 2017
- (51) Urbanowicz, Ryan J., and Jason H. Moore. "ExSTraCS 2.0: description and evaluation of a scalable learning classifier system." *Evolutionary Intelligence* (2015): 1-28.
- (52) Urbanowicz, Ryan J., and Jason H. Moore. "The application of michigan-style learning classifier systems to address genetic heterogeneity and epistasis in association studies." *Proceedings of the 12th annual conference on Genetic and evolutionary computation*. ACM, 2010.

References (5/5)

- (53) Urbanowicz, Ryan J., Ambrose Granizo-Mackenzie, and Jason H. Moore. "An analysis pipeline with statistical and visualization-guided knowledge discovery for michigan-style learning classifier systems." *Computational Intelligence Magazine, IEEE* 7.4 (2012): 35-45.
- (54) Urbanowicz, Ryan, Ambrose Granizo-Mackenzie, and Jason Moore. "Instance-linked attribute tracking and feedback for michigan-style supervised learning classifier systems." *Proceedings of the 14th annual conference on Genetic and evolutionary computation*. ACM, 2012.
- (55) Urbanowicz, Ryan J., Delaney Granizo-Mackenzie, and Jason H. Moore. "Using expert knowledge to guide covering and mutation in a michigan style learning classifier system to detect epistasis and heterogeneity." *PPSN XII*. Springer Berlin Heidelberg, 2012. 266-275.
- (56) Urbanowicz, R. J.; Bertasius, G. & Moore, J. H., "An Extended Michigan-Style Learning Classifier System for Flexible Supervised Learning, Classification, and Data Mining", PPSN XIII, Springer International Publishing, 2014, 211-221
- (57) Urbanowicz, R. J.; Lo, C.; Holmes, J. H. & Moore, J. H., "Attribute Tracking: Strategies Towards Improved Detection and Characterization of Complex Associations", GECCO 2018, ACM, 2018, 553-560
- (58) Urbanowicz, R. J. & Browne, W. N., "Introduction to Learning Classifier Systems", Springer Publishing Company, 2017
- (59) Wilson, Stewart W. "ZCS: A zeroth level classifier system." *Evolutionary computation* 2.1 (1994): 1-18.
- (60) Wilson, Stewart W. "Classifier fitness based on accuracy." *Evolutionary computation* 3.2 (1995): 149-175.
- (61) Wilson, Stewart W. "Get real! XCS with continuous-valued inputs." *Learning Classifier Systems*. Springer Berlin Heidelberg, 2000. 209-219.
- (62) Wilson, Stewart W. "Classifiers that approximate functions." *Natural Computing* 1.2-3 (2002): 211-234.
- (63) Wilson, S., "Three Architectures for Continuous Action", Learning Classifier Systems, Springer Berlin Heidelberg, 2007 , 4399 , 239-257
- (64) Wilson, S. W., "Classifiers that Approximate Functions", Natural Computing, Kluwer Academic Publishers, 2002, 1, 211-234