

# Course Agenda Introduction How LCSs map a problem Demo of 'Classifiers' How LCS learn a better map Demo of the main LCSs evolutionary cycle Organisation of a LCS Demo of different concepts within LCSs Applications of LCSs Demo of different types of LCSs Overview, Questions & Discussion

### Guides

- Will N. Browne's main area of research is Artificial Cognitive Systems. He has served as Co-track chair for Genetics-Based Machine Learning in GECCO 2011 and 2012 [plus organizing committee of the International Workshop on Learning Classifier Systems (IWLCS) from 2009-2010].
   Editor in chief for the Australasian Conference on Robotics and Automation 2012 and organized the Cognitive Robotics Intelligence and Control COGRIC, EPSRC (UK)/NSF (USA)
- Ryan Urbanowicz holds a Ph.D in genetics from Dartmouth College, and both a M.Eng. and B.Eng. in agricultural and biological engineering from Cornell University. His current research focuses on the development of machine learning strategies for feature selection, modeling, classification, and data mining in studies of common complex human disease. He served on the IWLCS organizing committee from 2010-2012 and is returning as an organizer from 2012-2014.

2

### Introduction This tutorial will introduce the concept of Learning Classifier Systems (LCS) User-friendly in order to allow Graduate students, EC researchers and Industry/business practitioners

who want to get up to speed with the field an easy path into using LCS to solve their complex problems.

### Introduction



- "LCS are a quagmire -
- a glorious, wondrous and inventing quagmire, but a quagmire nonetheless" D. Goldberg '92

### Not anymore!

- 30+ years research on LCS has clarified understanding, produced algorithmic descriptions, determined 'sweet spots' for parameters and delivered understandable 'out of the box' code
- This tutorial/book offers a boardwalk through the swamp

### Introduction



- Deliverables:
- This tutorial offers a user-friendly guide so that you will be able to proficiently implement LCS.
- It will be through explanation based on the
  - slides accompanying the book,
  - · examples supported by the Python code and
  - · insight/narrative from the two authors.
- Tutorial participants will have online access to the textbook, slides and code to work through the examples.

### Introduction



- Wondrous as:
- Learning Classifier Systems combine the global search of *Evolutionary Algorithms* with the local optimisation of *Reinforcement Learning* to address classification and regression problems.
- The knowledge extracted though interacting with data or embedded in an environment is human readable.
- 'Inventing' as LCS' flexible nature allows application to many domains with many types of feedback on solution progress.
- But 'swampy' as an LCS is not a one line algorithm with separable methods and easily tuned parameters.

6

### By The Way



- "Learning Classifier Systems" is an odd name!
- There are many Artificial Intelligence systems that learn to classifier (such as Decision Trees) that are not Learning Classifier Systems
- Learning Classifier System, abbreviated to LCS, refers to a singular/specific system, whereas Learning Classifier Systems (LCSs) refers to multiple systems or the field.
- Genetics-Based Machine Learning (GBML) is more accurate, but still not completely precise (e.g. Artificial Immune Systems are GBML, but not LCSs).
- Please accept the limitations in the name and let's explore the the concepts that underline LCSs.

### LCSs [History 1 of 1]:



- Learning Classifier Systems are one of the earliest artificial cognitive systems.
- The early work was ambitious and broad. This has led to many paths being taken to develop the concept over the next 30 years.
- Coupling this with the fact that replicating cognition in itself is a difficult problem has led to the field being affectionately termed 'a quagmire' and a lack of widespread adoption.
- Early 90s simplified and essential ideas, then understanding and adaption to real-world problems.

10

### **Assumptions for Al:**



12

- In order for artificial learning to occur data containing the patterns to learn is needed.
- This can be through recorded past experiences or interactive with current events.
- Learning is often in the harness of a cognitive system as input data, representation, reasoning, learning and output are needed to interact with an environment

If there are no clear patterns in the data, then LCSs will not learn.

### Why learn using AI:



- Learning is valued by humans as it enhances our abilities to solve problems and adapt to our environment.
- Much work in research fields, such as education, psychology and neuroscience, has been conducted into how humans learn.
- With the advent of computers, humans have been interested in seeing how artificial 'agents' could learn. Either learning to
  - solve problems of value that humans find difficult to solve
  - for the curiosity of how learning can be achieved.



### Worth of a Rule



An 'if ... then ... ' rule may be valid syntactically, but we need to verify its worth

- A valid rule can quite easily encode meaningless relationships and information.
- Interestingly, the majority of valid rules are likely to contain incorrect information
- In LCSs the worth of a rule is termed 'fitness', due to analogies of biological fitness.







the learning signal.



Querying

Data-mining

21

### **Environmental interaction Tasks** GECO GE Search Optimisation Supervised learning: The environment contains a teacher that (directly or indirectly) Modelling Scheduling provides the correct response for certain environmental states as a training signal for **Knowledge-Handling** Prediction Unsupervised learning: The learning system has an Visualisation Design internally defined teacher Routing with a prescribed goal that Learning does not need utility Adaptive-control Game-playing feedback of any kind. Reinforcement learning: The **Rule-Induction** environment does not directly indicate what Diagnosis the correct response should have been. Instead, it only provides reward or punishment to indicate the utility of actions that were actually taken by the system. Classification [Sutton and Barto 98] 19









### Competition

111						1		
110								
101							1	
100			1	1			1	
011					1	1	1	
010					-1	1		
001								
000								
	0	1	0	1	0	1	0	1
	0	0	1	1	0	0	1	1
	0	0	0	0	1	1	1	1

- Ideally, there would only be one unique and correct rule for each niche
- Number of rules would equal number of niches

No prior knowledge, so each rule must be learnt.

- LCSs allow multiple, slightly different rules per niche
- Multiple hypotheses are available to find the optimum rule
- Each rule 'covers', i.e. describes, its part of the search space.
- The rules within a niche compete to map the domain.

26

### Learning 'learning' has a very useful definition

"Learning is constructing or modifying representations of what is being experienced" Michalski et al., 86.

- LCSs need to experience the domain in order to learn.
  - · embodied in an actual robot or
  - · 'virtual' as in a software program receiving data
- Noise and dynamics within the data may impact on learning ability, but LCSs have shown robustness
- LCSs construct rules or modify existing rules in order to learn



### **Search space** A major factor in determining the difficulty/likelihood of success. Every valid rule can be thought of a candidate solution, i.e. rules that satisfy problem constraints and encoding Note, we do not yet know if it is a good solution to the problem or not. Each candidate solution is a member of the set of possible solutions. The space of all candidate solutions is termed the 'search space' [alternative names, such as feasible set, feasible region and solution space exist, but are rarely used in LCS literature]. The size of the search space is determined by both the encoding of the LCS itself and the problem itself. 29

### Representation



- LCSs can use multiple representation schemes.
   Suited to binary input or

Environment input must be encoded

- Suited to binary input of
  Suited to real-valued inputs and so forth...
- Consider representing the hours in a day in four bit binary
- $\boldsymbol{\diamond}$   $\hfill \hfill \hfill$
- The distance between similar genotypes (and their expressed phenotypes) in a search space is an important consideration when deciding upon the representation

30

### **Don't Care**

A condition that we don't care about is given the symbol '#'



32

### For example,

101:1- the Boolean states 'on off on' has action 'on'001:1- the Boolean states 'off off on' has action 'on'Can be encoded as

#01:1 - the Boolean states '. off on' has action 'on'

- The ternary alphabet in the Classifier matches binary input
- In many instances, # acts as an OR function on {0,1}

Redundancy, irrelevance and compactness



31

- \* A search space has a number of dimensions.
- \* A single condition encodes a single dimension (feature).
- If the problem has multiple dimensions (features) then the LCS will need the corresponding number of conditions.
- The division between the conditions is implicit within a classifier. For example, the condition string: '100110'
  - 1x6-bit number, 2x3-bit numbers or 6xBoolean state
- Not all conditions are useful in the map

<section-header><section-header><section-header><section-header><list-item><list-item><list-item><list-item><list-item><list-item><list-item><list-item><list-item><list-item>





![](_page_8_Figure_2.jpeg)

![](_page_8_Picture_3.jpeg)

### **Cover method**

![](_page_9_Picture_1.jpeg)

41

42

Verb 'to Cover', or noun 'Coverage'

Rule creation:

- Condition: Generalisation of the environmental instance
- Action: Known (supervised learning) or Random (reinforcement learning)

Environmental input: 101101 with corresponding action 1 with probability of generalisation (P#) of 0.33 Rule\_ 1##101:1 + statistics

### Match

![](_page_9_Picture_8.jpeg)

43

Special cases in the match method:

- \* No matching classifiers then invoke cover method
- Not sufficient matching classifiers then invoke cover method
- Not sufficient matching classifiers with all possible actions then invoke cover method (useful for a complete map)
- Partial match of a Classifier acceptable (e.g. 5/6 states)?
  - Low dimensional problems, this obscures class boundaries, so no.
  - May be necessary in Big Data

Match
Do any of the rules match the input?
Match method: for example, new input 111101 also action 1.
Specific: rule conditions exactly match the input 1----- matches Rule<sub>a</sub> 1##101:1
General: 'don't care' matches all input -1---- matches Rule<sub>a</sub> 1##101:1
Environmental input: 111101 with corresponding action 1 Matches the previously generated rule Rule<sub>a</sub> 1##101:1 + statistics
All matching rules placed in the match set [*M*] Explore vs. Exploit
Image: problem of the biggest problems in evolutionary computation
When to explore to learn new knowledge?
When to explore to learn new knowledge?
Annealing schemes must be set a priori without knowledge of the optimum scheme
Often just a fixed ratio!

![](_page_10_Figure_0.jpeg)

![](_page_10_Figure_1.jpeg)

![](_page_10_Figure_2.jpeg)

![](_page_11_Figure_0.jpeg)

![](_page_11_Figure_1.jpeg)

![](_page_11_Figure_2.jpeg)

![](_page_11_Picture_3.jpeg)

![](_page_12_Figure_0.jpeg)

![](_page_12_Figure_1.jpeg)

![](_page_12_Figure_2.jpeg)

![](_page_12_Figure_3.jpeg)

![](_page_13_Figure_0.jpeg)

![](_page_13_Figure_1.jpeg)

![](_page_13_Figure_2.jpeg)

![](_page_13_Figure_3.jpeg)

### **Rights + Wrongs**

![](_page_14_Picture_1.jpeg)

- Accuracy measures the inverse of predictive error
- What happens if you are always wrong?
- If you can predict this accurately, then you are 100% accurate!
- In some cases, e.g. MUX, this doubles the rule base

C	Α	Р	F
01#	1	100	100
00#	0	100	100
1#1	1	100	100
1#0	0	100	100
01#	0	0	100
00#	1	0	100
1#1	0	0	100
1#0	1	0	100

![](_page_14_Figure_7.jpeg)

![](_page_14_Picture_8.jpeg)

![](_page_14_Figure_9.jpeg)

Choic Euclidea	Choice of Encoding Euclidean and Hamming distances alter search space								
Integer 0 1 :1 2 :1 3 :1 4 :1 5 :1 6 :1 7 :1	Binary 000 001:1 010:2 011:1 100:3 101:1 110:2 111:1	Gray 000 001 :1 011 :1 010 :1 110 :1 101 :1 100 :1	Enumerated 0000000 0000001 :1 0000011 :1 0000111 :1 0001111 :1 0011111 :1 0111111 :1	d Encoding :Hamming distance					
How to e	encode the	e range: 0	ightarrow 3 and 0 $ ightarrow$	→ <b>4</b> ? <sub>68</sub>					

Alphabet	Rule	Match
ernary	1#1000 : 0	1#1000
ernary	11#00# : 0	11#00#
nteger	111000 u:0	1#1000
-	101000 l	
nteger	111001 u:0	11#00#
	110000 l	
Real	$1.0 \ 1.0 \ 1.0 \ 0.2 \ 0.3 \ 0.9 \ u \ : \ 0.0$	1#1000
	$0.5 \ 0.0 \ 0.7 \ 0.0 \ 0.0 \ 0.0 \ 1$	
Real	$1.0 \ 1.0 \ 1.0 \ 0.2 \ 0.3 \ 1.0 \ u \ : \ 0.0$	11#00#
	$0.5 \ 0.0 \ 0.0 \ 0.0 \ 0.0 \ 0.9 \ 1$	

### **Integer and Real Encodings**

![](_page_15_Picture_3.jpeg)

- Encoding suited to the environmental message
  - using upper and lower bounds:

Alphabet	Rule	Match
Integer	111001 u : 0	11#00#
	110000 l	
Real	$1.0 \ 1.0 \ 1.0 \ 0.2 \ 0.3 \ 0.9 \ u$ : $0.0$	1# <b>1</b> 00 <b>0</b>
	$0.5 \ 0.0 \ 0.7 \ 0.0 \ 0.0 \ 0.0 \ 1$	

[could use centre and spread, but this assumes a Gaussian distribution and recombination more difficult to implement]

- ♦ For each allele *a*,  $lb \le x \le ub$ , to give match.
- ♦ Could use '<' instead of '≤' but LCSs determine the correct bound automatically
  - e.g.  $0 \le x \le 5$  is equivalent to  $0 \le x \le 5.01$

![](_page_15_Picture_11.jpeg)

- ♦ Note that most bounds have a limit, e.g. WBC:  $0 \le a \le 10$ 
  - We decide to mutate the lower bound of a '#'  $0 \le x \le 10$
  - If decrease by 10% of range to -1, repair back to 0.
  - If increase by 10% of range to 1, do not repair as valid!

### Thus some alphabets have a specificity bias

![](_page_16_Figure_0.jpeg)

separate with this encoding, so use Hyperellipsoids

![](_page_16_Figure_2.jpeg)

![](_page_16_Figure_3.jpeg)

![](_page_16_Figure_4.jpeg)

### **Overgenerals**

[undesired inaccurate classifiers]

![](_page_17_Picture_2.jpeg)

- When additional reward offsets any additional penalty
- Strength-based fitness is more prone to overgenerals
- Accuracy-based fitness is less prediction orientated

Want 10011###1:1 get 10011####:1, where 10011###0:0

- Can occur in unbalanced datasets or
- where the error tolerance  $\varepsilon_0$  is set too high.

![](_page_17_Figure_9.jpeg)

### **Subsumption deletion**

![](_page_17_Picture_11.jpeg)

77

- In sparse environments over-specific rules can take over population Want 10011###1:1 get 10011#011:1, 100111111:1, ...
- Starvation of generals, so delete specific 'sub-copies'
- Need accurate rules first: how to set level of accuracy (often not 100%)
- GA or action set [A] subsumption?
- Effect of numerosity?

![](_page_17_Picture_18.jpeg)

### Set Pressure

![](_page_18_Picture_1.jpeg)

- Set pressure is related to the opportunity to breed,
- Does not occur in panmictic rule selection
- Need Niching through [M] or [A] rule discovery
- Class imbalance affects set pressure
- Set pressure is more effective when replacing 'weaker' rules
- Often panmictic deletion, thus one action can replace a different action
- To prevent an action type disappearing, relative fitness is used (rare rules have high relative fitness and so breed)
- Rules that occur in more sets have more opportunity to be selected from mating
- Favours general rules

![](_page_18_Figure_11.jpeg)

![](_page_18_Figure_12.jpeg)

![](_page_18_Picture_13.jpeg)

![](_page_19_Figure_0.jpeg)

![](_page_19_Figure_1.jpeg)

### Numerosity

Numerosity affects select and update procedures:

![](_page_19_Picture_4.jpeg)

- When calculating the fitness of an action in the select for effect procedure:
- The fitness used assumes all the classifiers in the population:

$$f = \frac{\sum F \times p \times n}{\sum F \times n}$$

Where n is the numerosity

- ♦ Macroclassifiers: all unique classifiers n ≥ 1
- Microclassifiers: all individual classifiers (n copies of macroclassifiers)
- Ratio of macroclassifiers to microclassifiers often used as a measure of training progress

![](_page_19_Picture_12.jpeg)

### XCS

![](_page_20_Picture_1.jpeg)

Arguably the most adopted LCS:

- Wilson 1995: Evolutionary Computation v3n2. pp1-44
- Not an acronym!
- Although rumoured to be eXtended Classifier System
- Extends basic ZCS
- Niche based fitness to add set pressure
- Accuracy based LCS
- Complete and accurate mapping from inputs and actions to pay off predictions
- Maximally general subject to an accuracy criterion (aided by subsumption)

93

94

### **Classifier parameters**

Symbol	Meaning
р	Prediction p Estimates (keeps average of) the payoff expected if the classifier matches and its action is taken by the system
ε	Prediction error Estimates the errors made in the predictions
F	Fitness Denotes the classifier's fitness
exp	Experience Counts the number of times since its creation that the classifier has belong to an action set
ts	Time stamp Denotes the time-step of the last occurrence of a GA in an action set to which this classifier belonged
as	Action set size Estimates the average size of the action sets this classifier has belonged to
n	Numerosity Reflects the number of micro-classifiers (ordinary classifiers) this classifier – which is technically called a macroclassifier - represents
	9

## Complete map Should LCS discover: The most optimum action in a niche The predicted payoff for all actions in a niche X x A => P (cf Q-Learning) The danger with optimum action only: If a suboptimal rule is converged upon ... difficult to discover and switch policy (CF path habits) The problem with predicting all actions: Memory and time intensive Identifies and keeps consistently incorrect action (100% accurate prediction) rules

Harder to interpret rule base

Learning parameters Maximum size of the population (In micro-classifiers, N is the number sum of the classifier numerosities) β Learning rate for p,  $\epsilon$ , f, and as Used in calculating the fitness of a classifier α, ε<sub>0</sub>, ν Discount factor used (in multi-step problems) in updating classifier predictions θω GA threshold The GA is applied when the average time since the last GA in the lest is greater than the threshold Probability of applying crossover in GA χ μ Probability of mutating an allele in the offspring θ<sub>de</sub> Deletion threshold If the experience of a classifier is greater than  $\theta_{del}$ , its fitness may be considered in its probability of deletion δ Mean fitness in [P] below which the fitness of a classifier may be considered in its probability of deletion Subsumption threshold – the experience of a classifier must be greater than  $\theta_{sub}$  in order to  $\theta_{sub}$ be able to subsume another classifier 98

Symbol	Meaning
P <sub>#</sub>	Probability of using a # in one attribute in C when covering
$p_I \epsilon_I F_I$	Initial Values in new classifiers
p <sub>explr</sub>	Probability during action selection of choosing the action uniform randomly
θ <sub>mna</sub>	Minimal number of actions that must be present in a match set [M] or else covering will occur
doGASumputi	on is a Roolean parameter that specifies if offenring are to be tested for possible logical
subsumption	on la a bourean parameter that specifies in onspiring are to be tested for possible logical by parents.
doActionSetSi classifiers.	by parents.
doActionSetSi classifiers.	ubsumption is a boolean parameter that specifies if action sets are to be tested for subsuming
doActionSetSu classifiers.	by parents.
doActionSetSu classifiers.	by parents.
doActionSetSi classifiers.	up a a boocan parameter that specifies if action sets are to be tested for subsuming
doActionSetSi classifiers.	ubscan parallelet that specifies if onspiring are to be tested for possible logical parents.

<b>Results Interpretation</b>	n
-------------------------------	---

No.	Condition	Action	n	F	е	р	exp
1	0010##0################################	0	24	0.81	0	1000	455
2	00001##################################	0	23	0.92	0	0	7281
3	0010##0################################	1	22	0.79	0	0	7007
4	0101#####1#############################	1	19	0.62	0	1000	7542
5	0111#######0###########################	1	18	0.66	0	0	7358
6	0110######1############################	1	17	0.69	0	1000	7324
7	0111#######1###########################	0	15	0.53	0	0	6013
8	1010#########1######	1	14	0.64	0	1000	3841
9	00000##################################	0	11	0.54	0	1000	74
10	00#01#1################################	1	8	0.22	0	1000	1131
11	011#######00#########	1	8	0.23	0	0	1810
12	11#0###########0#0#	0	8	0.27	0	1000	2483
13	0#000###0############	0	8	0.29	0	1000	6435
14	00#1#1#1###############################	0	8	0.23	0	0	7170
15	000#00##################	1	7	0.21	0	0	193

GECO

Note: 2, 3, 5, 7, 11, 14, 15 incorrect, but accurate, so fit

![](_page_21_Figure_4.jpeg)

![](_page_21_Figure_5.jpeg)

![](_page_22_Figure_0.jpeg)

### **Cognitive Systems definition**

![](_page_22_Picture_2.jpeg)

Current robots are poor cognitive systems. Need to improve devices that we use every day and investigate medical benefits.

"Cognitive systems are natural or artificial information processing systems, including those responsible for perception, learning, reasoning and decision-making and for communication and action"

DTI Foresight initiative.

### **Cognitive LCS**

![](_page_22_Picture_7.jpeg)

Induction: Processes of Inference, Learning and Discovery (Computational Models of Cognition and Perception)

### J. H. Holland, K. J. Holyoak, R. E. Nisbett and P. Thagard

Two psychologists, a computer scientist, and a philosopher have collaborated to present a framework for understanding processes of inductive reasoning and learning in organisms and machines. Theirs is the first major effort to bring the ideas of several disciplines to bear on a subject that has been a topic of investigation since the time of Socrates. The result is an integrated account that treats problem solving and induction in terms of rulebased mental models.

### MIT Press (1 Jan 1986) ISBN: 0262081601

### Anticipatory LCS

- Anticipations influence cognitive systems and illustrates the use of anticipations for
- 1. Faster reactivity
- 2. Adaptive behavior beyond reinforcement learning
- 3. Attentional mechanisms
- 4. Simulation of other agents
- 5. The implementation of a motivational module.
- A particular evolutionary model learning mechanism, a combination of
  - a directed specializing mechanism and
  - a genetic generalizing mechanism.
- Experiments show that anticipatory adaptive behavior can be simulated by exploiting the evolving anticipatory model for even faster model learning, planning applications, and adaptive behavior beyond reinforcement learning.

![](_page_23_Figure_0.jpeg)

Abstr	action	GECCO
Connect4 simplified	4 is a noiseless domain so th 1:	ne fitness update can be
-> p ← p ← p ←	p + 2 win p draw p - 2 lose	Min 0 Max 100
k ← k ←	k + 2 if prediction is correct k – 2 otherwise	t,

	Abstracted Rules					
e.g.	'if side guide setting < width, then poor quality product					
	Abstraction checks for patterns in the base rules and crates and abstracted rules for each discovered pattern					
Base ru	les					
e.g. if si if side-g	de-guide-setting = 80, width = 82 then poor quality product juide-setting = 79, width = 80 then poor quality product					
e.g. if si if side-g	de-guide-setting = 80, width = 82 then poor quality product uide-setting = 79, width = 80 then poor quality product Learning system					
e.g. if si if side-c Raw D e.g. Fe	de-guide-setting = 80, width = 82 then poor quality product uide-setting = 79, width = 80 then poor quality product Learning system ata atures side-guide setting', ' width' : 'product quality'					
Raw D e.g. ff side-c Raw D e.g. Fe	de-guide-setting = 80, width = 82 then poor quality product uide-setting = 79, width = 80 then poor quality product Learning system ata atures side-guide setting', ' width' : 'product quality' 81 : poor					
Raw D e.g. Fe e.g. Fe 78 79	de-guide-setting = 80, width = 82 then poor quality product uide-setting = 79, width = 80 then poor quality product Learning system ata atures side-guide setting', ' width' : 'product quality' 81 : poor 80 : poor					
Raw D e.g. fe e.g. Fe 78 79 78	de-guide-setting = 80, width = 82 then poor quality product uide-setting = 79, width = 80 then poor quality product Learning system ata atures side-guide setting', ' width' : 'product quality' 81 : poor 80 : poor 76 : good					

![](_page_23_Figure_3.jpeg)

![](_page_24_Figure_0.jpeg)

![](_page_24_Figure_1.jpeg)

![](_page_24_Picture_2.jpeg)

![](_page_24_Picture_3.jpeg)

![](_page_25_Figure_0.jpeg)

Rule Compaction Algorithm – Testing Accuracy Based Retain a non-redundant, maximally general subset of the rule population.

Custom Fitness Calculation Based a rule's accuracy, generality, with and strength (age restricted)

Visualization Strategies – Interpreting the Black Box

![](_page_25_Figure_4.jpeg)

![](_page_25_Figure_5.jpeg)

# <section-header>

	ndivi	ndividual Attributes						Pairs				
	Attribute	SpS	p-value	AWSpS	p-value		Attrib	oute Pairs	CoS	p-value		
	X0	10885	$0.001^{*}$	7589.49	0.001*		X0	X1	8060	0.001*		
.50 .50	X1	11359	0.001*	7936.43	0.001*		X2	X3	7373	0.001*		
	X2	10569	0.001*	7369.84	0.001*		X1	X2	4223	0.001*		
	X3	10150	0.001*	7114.25	0.001*		X0	X2	4079	0.001*		
X0 X1 X2 X3	X4	3863	0.999	2482.56	0.888		X1	X3	3974	0.001*		
	X6	5240	0.737	2090.05	0.18		X0	X3	3829	0.001*		
Attributes: 20	X7	5484	0.915	3647.67	0.336		X1	X11	3621	0.001*		
Predictive: 1	X8	4429	0.95	2927.85	0.482		X1	X7	3574	0.001*		
Nep Dredictive: 16	X9	5334	0.985	3569.25	0.484		X0	X11	3540	0.001*		
Non-Predictive: 16	X10	5907	0.414	3948.81	0.04*		X2	X10	3485	0.001*		
Heritability = 0.4	X11	5725	0.414	3933.61	0.037*		X0	X7	3462	0.001*		
MAF = 0.2	X12	5273	1	3518.87	0.761		X3	X10	3392	0.001*		
Sample Size = 1600	X13	4443	1	2854.43	0.996		X1	X18	3379	0.001*		
	X14	3709	1	2391.91	0.978		X0	X18	3264	0.001*		
	X15	5108	0.916	3425.64	0.355		X1	X15	3255	0.001*		
Testing Accuracy = 0.70 (p	X16	3613	1	2299.56	1		xo	X15	3035	0.001*		
= 0.001)	X17	3639	1	2302.01	0.992		x2	X12	3020	0.005*		
	X18	5933	0.629	4004.96	0.085		N2	V19	2016	0.001*		
	X19	4591	1	2940.28	0.999		A2	A18	3016	0.001*		

![](_page_26_Figure_0.jpeg)

![](_page_26_Figure_1.jpeg)

![](_page_26_Figure_2.jpeg)

![](_page_26_Figure_3.jpeg)

![](_page_27_Figure_0.jpeg)

### Bladder Cancer Analysis: 10 Attributes Testing Accuracy = 0.60 (p = 0.001)Individual Attributes Pairs Attribute SpS p-value AWSpS p-value Attribute Pairs CoS & p-value XPD.751 5686 0.001\* 4001.86 0.001\* XPD.751 XPD.312 3367 0.001\* XPD.312 5077 0.007\* $3550.64 \quad 0.001^*$ 2757 0.001\* XPD.751 pack.yr 4827 0.037\* 3383.68 0.006\* pack.vr XPD.751 XRCC1.194 2530 $0.001^{*}$ XRCC1.194 4206 0.461 2818.81 0.179 age.504151 0.721 2800.63 0.308 XPD.312 pack.yr 2375 $0.006^{*}$ 4048 0.745 male $2752.61 \quad 0.358$ XPD.751 age.502345 $0.016^{*}$ XRCC1.399 3797 0.729 2595.78 0.427 XRCC1.194 2120 $0.04^{*}$ pack.yr APE1 3524 0.891 2418.03 0.667 XRCC1.194 XPD.312 2105 0.046\* XRCC3 3172 0.989 2166.090.91XPC.PAT 2975 1.0 1994.75 0.98

![](_page_27_Figure_2.jpeg)

![](_page_27_Picture_3.jpeg)

![](_page_28_Figure_0.jpeg)

![](_page_28_Figure_1.jpeg)

![](_page_28_Figure_2.jpeg)

![](_page_28_Figure_3.jpeg)

![](_page_29_Figure_0.jpeg)

### LCS Summary

![](_page_29_Picture_2.jpeg)

138

139

- Rule-based, multifaceted, machine learning algorithms
- Global search and learning through evolution mechanism
- Local search and adaption through reinforcement learning techniques – competition with cooperation
- Multitude of flexible implementations and representations
- Practical applications as now paths through the swamp.

### **LCSs Applications**

 Reinforcement Learning Problems
 Find an optimal behavioral policy represented a compact set of rules.

![](_page_29_Picture_10.jpeg)

Function Approximation Problems Find an accurate function approximation represented by a partially overlapping set of approximation rules.

Classification / Data Mining Problems Find a compact set of rules that classify all problem instances with maximal accuracy.

![](_page_29_Picture_13.jpeg)

### References

- Bacardit, J., and Butz, M. V. "Data Mining in Learning Classifier Systems: Comparing XCS with Gassist", in Advances at the frontier of Learning Classifier Systems 276 - 290, 2007
- Booker, L. B., "Triggered Rule Discovery in Classifier Systems" Proc. 3rd Int. Conf. on Genetic Algorithms, page 265-274, 1989
- Butz, M., and Wilson, S. W. "An algorithmic description of XCS." in Soft Computing: a fusion of foundations, methodologies and applications, 162-170, 2002
- Drugowitsch, J., "Design and Analysis of Learning Classifier Systems - A Probabilistic Approach. in *Studies in Computational Intelligence*, Springer 1-239, 2008
- Goldberg, D. E., et al. "What Makes a Problem Hard for a Classifier System?" First International Workshop on Learning Classifier Systems (IWLCS92), NASA Johnson Space 1992
- Kovacs, T., "Learning classifier systems resources" in *Soft Computing* 6(3-4): 240-243, 2002
- Michalski, R. S., et al., "Machine learning: A multistrategy approach" Morgan Kaufmann, 1986
- Sutton, R. S., and Barto, A. G. "Reinforcement learning: An introduction". MA: MIT Press, 1998.