An Algorithmic Description of XCS *

Martin V. Butz[1] and Stewart W. Wilson[2]

[1] Institute for Psychology III & Department of Computer Science
University of Würzburg, Germany
butz@psychologie.uni-wuerzburg.de
[2] University of Illinois at Urbana-Champaign
Prediction Dynamics, Concord, MA 01742
wilson@prediction-dynamics.com

Abstract. A concise description of the XCS classifier system's parameters, structures, and algorithms is presented as an aid to research. The algorithms are written in modularly structured pseudo code with accompanying explanations.

1 Introduction

XCS is a recently developed learning classifier system (LCS) that differs in several ways from more traditional LCSs. In XCS, classifier fitness is based on the *accuracy* of a classifier's payoff prediction instead of the prediction itself. Second, the genetic algorithm (GA) takes place in the action sets instead of the population as a whole. Finally, unlike the traditional LCS, XCS has no message list and so is only suitable for learning in Markov environments (XCS extensions using an internal-state register have shown promise in non-Markov environments).

XCS's fitness definition and GA locus together result in a strong tendency for the system to evolve accurate, maximally general classifiers that efficiently cover the state-action space of the problem and allow the system's 'knowledge' to be readily seen. As a result of these properties, there has been considerable interest in further investigation and potential extension of XCS and its principles. We therefore thought it would be useful to provide a basic algorithmic description of XCS, both as a core definition of the system and as a common framework from which new variants and research directions could spring.

We first present XCS's relation to the problem environment, followed by the system's structures and parameters. The rest of the paper consists of a top-down modular description of the XCS algorithm, written in pseudo-code accompanied by explanatory notes. We hope the result will be useful, and we encourage researchers to give us feedback regarding potential problems and clarifications. This document should definitely be read in conjunction with some of the basic XCS literature, for example [Wil95], [Kov97], and [Wil98]. Additional papers on XCS and other LCSs, together with a complete LCS bibliography, are found in [LSW00].

* Revised version of [BW01]

2 Environment Interaction, Structures, and Parameters

2.1 Interaction with the Environment

In keeping with the typical LCS model, the environment provides as input to the system a series of sensory situations $\sigma(t) \in \{0,1\}^L$, where L is the number of bits in each situation. In response, the system executes actions $\alpha(t) \in \{a_1, ..., a_n\}$ upon the environment. Each action results in a scalar reward $\rho(t)$ (possibly zero). The interaction is divided into *problems*, which may be either single-step or multi-step. A flag *eop* indicates the end of a problem. While $\sigma(t)$ and $\alpha(t)$ are interactions with the environment itself, the reward $\rho(t)$ and the flag are normally provided by another component which, following [DC98], we term the *reinforcement program rp*. The reinforcement program determines the reward according to the current environmental input and the action that was executed. The separation of environment and reinforcement components is useful and natural because reinforcement is often not an inherent aspect of the environment, but may be due, e.g., to a trainer or to the system's own priorities. Figure 1 illustrates the interaction of the environment and the reinforcement program with XCS.

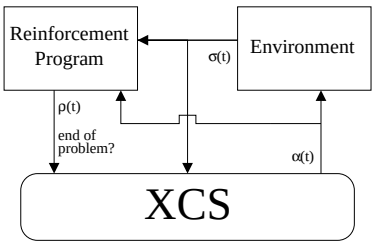

Fig. 1. XCS interacts with an environment and a reinforcement program.

In a single-step problem such as the Boolean multiplexer, the successive situations are not related to each other. After execution of an action the *rp* provides the appropriate reward ρ and signals with the flag *eop* that the problem has ended. In a multi-step problem such as a maze, the successive situations are related to each other. Reward is only provided in certain situations (e.g. the food position(s) in mazes). The point at which a problem ends in a multi-step environment must be defined by the *rp* according to the task that is to be solved.

2.2 A Classifier in XCS

XCS keeps a population of classifiers which represent its knowledge about the problem. Each classifier is a condition–action–prediction rule having the following parts:

- The condition $C \in \{0, 1, \#\}^L$ specifies the input states (sensory situations) in which the classifier can be applied (matches).
- The action $A \in \{a_1, ..., a_n\}$ specifies the action (possibly a classification) that the classifier proposes.
- The prediction p estimates (keeps an average of) the payoff expected if the classifier matches and its action is taken by the system.

Moreover, each classifier keeps certain additional parameters:

- The prediction error ϵ estimates the errors made in the predictions.
- The fitness F denotes the classifier's fitness.
- The experience exp counts the number of times since its creation that the classifier has belonged to an action set.
- The time stamp ts denotes the time-step of the last occurrence of a GA in an action set to which this classifier belonged.
- The action set size as estimates the average size of the action sets this classifier has belonged to.
- The numerosity n reflects the number of micro-classifiers (ordinary classifiers) this classifier—which is technically called a *macroclassifier*—represents.

To refer to one of the attributes of a classifier cl we use the dot notation $cl.x$ where x can be any of the above attributes (i.e. $x \in \{C, A, p, \epsilon, F, exp, ts, as, n\}$).

Important in our notation is the term *payoff*. Due to the Q-learning-like reinforcement learning in XCS, payoff does not refer solely to the expected reward ρ but is a combination of ρ and the payoff prediction of the best possible action in the next state. However, in the case of a single-step problem payoff reduces to the reward produced by the proposed action.

Note that ϵ measures the error of the predictions in units of payoff. The same is true of the parameter ϵ_0 that is defined in section 2.4.

Classifiers in XCS are *macroclassifiers*, i.e., each classifier represents n traditional or *micro-* classifiers having identical conditions and actions. Algorithms covering creation, deletion, and adjustment of the numerosity of macroclassifiers are given in section 3.10, 3.11, and 3.12. In XCS, macroclassifiers are always handled *as though* they consist of n micro-classifiers.

2.3 The Different Sets

There are four different sets that need to be considered in XCS.

- The population $[P]$ consists of all classifiers that exist in XCS at any time t.
- The match set $[M]$ is formed out of the current $[P]$. It includes all classifiers that match the current situation $\sigma(t)$.
- The action set $[A]$ is formed out of the current $[M]$. It includes all classifiers of $[M]$ that propose the executed action.
- The previous action set $[A]_{-1}$ is the action set that was active in the last execution cycle.

2.4 Learning Parameters in XCS

In order to control the learning process in XCS the following parameters are used:

- N specifies the maximum size of the population (in micro-classifiers, i.e., N is the sum of the classifier numerosities).
- β is the learning rate for p, ϵ, f, and as.
- α, ϵ_0, and ν are used in calculating the fitness of a classifier.
- γ is the discount factor used—in multi-step problems—in updating classifier predictions.
- θ_{GA} is the GA threshold. The GA is applied in a set when the average time since the last GA in the set is greater than θ_{GA}.
- χ is the probability of applying crossover in the GA.
- μ specifies the probability of mutating an allele in the offspring.
- θ_{del} is the deletion threshold. If the experience of a classifier is greater than θ_{del}, its fitness may be considered in its probability of deletion.
- δ specifies the fraction of the mean fitness in $[P]$ below which the fitness of a classifier may be considered in its probability of deletion.
- θ_{sub} is the subsumption threshold. The experience of a classifier must be greater than θ_{sub} in order to be able to subsume another classifier.
- $P_{\#}$ is the probability of using a # in one attribute in C when covering.
- p_I, ϵ_I, and F_I are used as initial values in new classifiers.
- p_{explr} specifies the probability during action selection of choosing the action uniform randomly.
- θ_{mna} specifies the minimal number of actions that must be present in a match set $[M]$, or else covering will occur.
- *doGASubsumption* is a Boolean parameter that specifies if offspring are to be tested for possible logical subsumption by parents.
- *doActionSetSubsumption* is a Boolean parameter that specifies if action sets are to be tested for subsuming classifiers.

2.5 Commonly Used Parameter Settings

For parameter settings, it is best to check the literature for a similar experiment. In some cases, the following suggestions could be taken as starting points. The population size, N, should be large enough so that, starting from an empty population, covering occurs only at the very beginning of a run. The learning rate, β, could be in the range 0.1-0.2. The parameter α is normally 0.1. The parameter ϵ_0 is the error below which classifiers are considered to have equal accuracy; a typical value would be about one percent of the maximum value of ρ, e.g., 10 if the maximum value is 1000. The power parameter ν is typically 5. The discount factor γ has been 0.71 in many problems in the literature, but larger or smaller values could certainly work, depending on the environment. The threshold θ_{GA} is often in the range 25-50. Crossover probabilities χ in the range 0.5-1.0 have been used. Mutation probabilities μ in the range 0.01-0.05 are often used.

The deletion threshold θ_{del} could be about 20. δ is often taken to be 0.1. The subsumption threshold θ_{sub} could be about 20, though larger values (more experience) are important in some problems. $P_{\#}$ could be around 0.33. Larger values reduce the need for covering, but may make it harder to evolve accurate classifiers. The initialization parameters p_I, ϵ_I, and F_I should be taken very small—essentially zero. The exploration probability p_{explr} could be 0.5, but this depends on the type of experiment contemplated. To cause covering to provide classifiers for every action, choose θ_{mna} equal to the number of available actions.

The setting of *doGASubsumption* and *doActionSetSubsumption* depends on the problem. In general, subsumption is used to eliminate classifiers that clearly add nothing to the system's decision capability. In the case of GA subsumption, these are offspring classifiers whose conditions are logically subsumed by a parent's condition, given that the parent is both accurate and sufficiently experienced. In the case of action set subsumption, a general, accurate, and experienced classifier in an action set eliminates classifiers in the set that its condition logically subsumes. The two subsumption methods are independent and different in their effects. Broadly, action set subsumption appears to be 'stronger': it causes greater 'condensation' of the population. Subsumption is useful in problems where there is a well-defined underlying target function, such as the Boolean multiplexer.

Note, however, that subsumption is *not* necessary for XCS to form accurate maximal generalizations (see [Wil95] for the basic XCS *generalization hypothesis*). Subsumption tends to result in accurate classifiers that are as formally general as possible without being contradicted by any actually occurring input. Without subsumption, the system produces these as well as accurate but more-specific classifiers that match the same inputs as the general ones do [Wil98]. Subsumption produces smaller final populations, but because the more-specific classifiers are not present, the system is more vulnerable to environmental changes—it does not have the more-specific classifiers to "fall back on".

3 An Algorithmic Description of XCS

This section presents the algorithms used in XCS. The description starts from the top level. First, the overall execution cycle is described. The following subsections specify the single parts in more detail. Each specified sub-procedure in this description is written in capital letters. When referring to a module other than XCS (e.g. the environment *env*), a colon is used. However, due to the diverse modifications and extensions of XCS published during the last several years, we first clarify to which 'XCS' we are referring in the subsequent description.

3.1 Which XCS?

Since the introduction of XCS in 1995 several additions and modifications have been reported. Some of them did not result in a change of the name 'XCS', but

they were an important step in increasing the robustness of the system. This section describes which changes of the basic XCS [Wil95] are included as well as which ones are omitted herein.

The update of the classifier parameters is modified in several ways from the original publication. First, the order of the update is changed but can easily be switched back to the original one. Second, the calculation of the accuracy measure κ is changed from the original exponential function to a power law function as used in [Wil00]. Finally, the MAM ("moyenne adaptive modifiée") technique is not used in the fitness update of a classifier. This method results in a stronger robustness against inaccurate classifiers.

Moreover, several changes in the discovery component are considered herein. First, the covering criterion is changed. While the original criterion had to consider the mean prediction in the population, the approach herein simply assures that a certain number of actions is present in each match set. The change increases the speed of the program while assuring the original properties of the mechanism. Second, the GA is applied in the action set and subsumption deletion is used as published in [Wil98]. However, the subsumption method is further detailed distinguishing between *action set subsumption* and *GA subsumption*. Next, mutation is modified describing a pure niche mutation. While this mutation style promises a faster convergence later in the run, an unrestricted mutation can help early in the run to solve harder problems. Finally, the deletion method combines the two original methods as proposed in [Kov99].

Despite their importance, we decided not to include Lanzi's modifications of *specify* ([Lan97], [Lan99a]) nor the enhancement for stochastic environments (XCSμ, [Lan99d]). Moreover, no enhancements of the representation of the classifier conditions are considered, e.g., as published in [Lan99b], introducing a messy coding, and in [Lan99c], introducing s-expressions. We also did not incorporate any sort of memory mechanism as recently investigated in detail (see e.g. [LW00]).

Thus, the XCS classifier system described herein includes the most important modifications while staying close to the original work. The modifications are further discussed in the algorithmic description itself.

3.2 Initialization

When XCS is started, the modules must first of all be initialized. The parameters in the environment must be set and, e.g., a maze must be read in. Also, the reinforcement program *rp* must be initialized. Finally, XCS itself must be initialized. Apart from the parameter settings and the start of the time-step counter referred to as `actual time` t, the population $[P]$ needs to be initialized. The population $[P]$ can either be left empty or can be filled with the maximal number of classifiers N, generating each classifier with a random condition and action and the initial parameters. The two methods differ only slightly in their effect on performance. Thus the simpler way of leaving the population empty in the beginning is commonly used. After the initialization, the main loop is called.

XCS:

```
1 initialize environment env
2 initialize reinforcement program rp
3 initialize XCS
4 RUN EXPERIMENT
```

3.3 The Main Loop

In the main loop *RUN EXPERIMENT*, the current situation is first sensed (received as input). Second, the match set is formed from all classifiers that match the situation. Third, the prediction array PA is formed based on the classifiers in the match set. PA predicts for each possible action a_i the resulting payoff. Based on PA, one action is chosen for execution and the action set $[A]$ is formed, which includes all classifiers of $[M]$ that propose the chosen action. Next, the winning action is executed. Then the previous action set $[A]_{-1}$ (if this is a multi-step problem and there is a previous action set) is modified using the Q-learning-like payoff quantity `P` which is a combination of the previous reward ρ_{-1} and the largest action prediction in the prediction array PA. Moreover, the GA may be applied to $[A]_{-1}$. If a problem ends on the current time-step (single-step problem or last step of a multi-step problem), $[A]$ is modified according to the current reward ρ and the GA may be applied to $[A]$. The main loop is executed as long as the termination criterion is not met. A termination criterion is, e.g., a certain number of trials or a 100% performance level.

The main loop specifies many sub-procedures essential for learning in XCS. Some of the procedures are more or less trivial while others are complex and themselves call other sub-procedures. The following sections describe all procedures specified in the main loop, covering all relevant processes. Each of them tries to specify the general idea and the overall process and then gives a more detailed description of single parts in successive paragraphs.

3.4 Formation of the Match Set

The *GENERATE MATCH SET* procedure gets as input the current population $[P]$ and the current situation σ. Although the procedure sounds trivial, it has within it a covering process. Covering is called when the number of different actions represented by matching classifiers is less than the parameter θ_{mna}. Thus, *GENERATE MATCH SET* first looks for the classifiers in $[P]$ that match σ and next, checks if covering is required. A classifier generated by covering can be directly added to the population since it must differ from all current classifiers. Note that the while loop at step 2 is executed as long as covering is required.

In the following paragraphs we will describe the sub-procedures included in the *GENERATE MATCH SET* algorithm. The sub-procedure *DELETE FROM POPULATION* however can be found in section 3.11.

RUN EXPERIMENT():

```
 1 ρ_{-1} ← 0
 2 do{
 3    σ ← env: get situation
 4    [M] ← GENERATE MATCH SET out of [P] using σ
 5    PA ← GENERATE PREDICTION ARRAY out of [M]
 6    act ← SELECT ACTION according to PA
 7    [A] ← GENERATE ACTION SET out of [M] according to act
 8    env: execute action act
 9    ρ ← rp: get reward
10    if([A]_{-1} is not empty)
11       P ← ρ_{-1} + γ * max(PA)
12       UPDATE SET [A]_{-1} using P possibly deleting in [P]
13       RUN GA in [A]_{-1} considering σ_{-1} inserting and
            possibly deleting in [P]
14    if(rp: eop)
15       P ← ρ
16       UPDATE SET [A] using P possibly deleting in [P]
17       RUN GA in [A] considering σ inserting and
            possibly deleting in [P]
18       empty [A]_{-1}
19    else
20       [A]_{-1} ← [A]
21       ρ_{-1} ← ρ
22       σ_{-1} ← σ
23 }while(termination criteria are not met)
```

GENERATE MATCH SET([P], σ):

```
 1 initialize empty set [M]
 2 while([M] is empty)
 3    for each classifier cl in [P]
 4       if(DOES MATCH classifier cl in situation σ)
 5          add classifier cl to set [M]
 6    if(the number of different actions in [M] < θ_{mna})
 7       cl_c ← GENERATE COVERING CLASSIFIER considering [M] and σ
 8       add classifier cl_c to set [P]
 9       DELETE FROM POPULATION [P]
10       empty [M]
11 return [M]
```

Classifier Matching The matching procedure is that commonly used in LCSs. A 'don't care'-symbol # in C matches any symbol in the corresponding position of σ. A 'care' or non-# symbol only matches with the exact same symbol at that position. The basic XCS relies on binary coding and thus the care symbols are $\in \{0, 1\}$. Recently, Lanzi ([Lan99b], [Lan99c]) introduced an extension of

the conditions in XCS involving messy coding and Lisp s-expressions. Wilson [Wil00] introduced the first XCS extension that is able to handle real coded inputs. However, here we will consider only the basic kind of classifier condition.

The *DOES MATCH* procedure checks each component in the classifier's condition C. If a component is specified (i.e. is not a don't care symbol), it is compared with the corresponding attribute in the current situation σ. Only if all comparisons hold does the classifier match σ and the procedure return *true*.

DOES MATCH(cl, σ):
```
1 for each attribute x in cl.C
2    if(x ≠ # and x ≠ the corresponding attribute in σ)
3       return false
4 return true
```

Covering Covering occurs if the number of actions present in $[M]$ is $< \theta_{mna}$. A classifier is created whose condition matches $\sigma(t)$ and contains don't cares with probability $P_{\#}$. The classifier's action is chosen randomly from among those not present in $[M]$.

GENERATE COVERING CLASSIFIER($[M]$, σ):
```
 1 initialize classifier cl
 2 initialize condition cl.C with the length of σ
 3 for each attribute x in cl.C
 4    if(RandomNumber[0,1) < P#)
 5       x ← #
 6    else
 7       x ← the corresponding attribute in σ
 8 cl.A ← random action not present in [M]
 9 cl.p ← pI
10 cl.ε ← εI
11 cl.F ← FI
12 cl.exp ← 0
13 cl.ts ← actual time t
14 cl.as ← 1
15 cl.n ← 1
16 return cl
```

3.5 The Prediction Array

Given an input, XCS makes a "best guess" prediction of the payoff to be expected for each possible action. These *system predictions* are stored in an array called the Prediction Array PA. The system prediction for an action is a fitness-weighted average of the predictions of all classifiers in $[M]$ that advocate that

action. If no classifiers in $[M]$ advocate a certain action, its system prediction is not defined, symbolized by *null*.

The *GENERATE PREDICTION ARRAY* procedure considers each classifier in $[M]$ and adds its prediction multiplied by its fitness to the prediction value total for that action. The total for each action is then divided by the sum of the fitnesses for that action to yield the system prediction.

GENERATE PREDICTION ARRAY($[M]$):

```
initialize prediction array PA to all null
initialize fitness sum array FSA to all 0.0
for each classifier cl in [M]
   if(PA[cl.A] = null)
      PA[cl.A] ← cl.p * cl.F
   else
      PA[cl.A] ← PA[cl.A] + cl.p * cl.F
   FSA[cl.A] ← FSA[cl.A] + cl.F
for each possible action A
   if(FSA[A] is not zero)
      PA[A] ← PA[A] / FSA[A]
return PA
```

3.6 Choosing an Action

XCS does not prescribe any particular action-selection method, and any of a great variety can be employed. For example, actions may be selected randomly, independent of the system predictions, or the selection may be based on those predictions—using, e.g., roulette-wheel selection or simply picking the action with the highest system prediction.

In our *SELECT ACTION* procedure we illustrate a combination of *pure exploration*—choosing the action randomly—and *pure exploitation*—choosing the best one. Lanzi [Lan99a] published first experiments with so-called *biased exploration* where pure exploration is chosen with a probability p_{explr} and otherwise pure exploitation is chosen. This action selection method is known as *ϵ-greedy* selection in the reinforcement learning literature [SB98] where the ϵ has identical meaning to our p_{explr} parameter. As an aside, it appears better to perform the GA only on exploration steps, especially if most steps are exploitation.

SELECT ACTION(PA):

```
if(RandomNumber[0,1) < p_explr)
   //Do pure exploration here
   return a randomly chosen action from those not null in PA
else
   //Do pure exploitation here
   return the best action in PA
```

3.7 Formation of the Action Set

After the match set $[M]$ is formed and an action is chosen for execution, the *GENERATE ACTION SET* procedure forms the action set out of the match set. It includes all classifiers that propose the chosen action for execution.

GENERATE ACTION SET($[M]$, act):
```
1 initialize empty set [A]
2 for each classifier cl in [M]
3    if(cl.A = act)
4       add classifier cl to set [A]
```

3.8 Updating Classifier Parameters

Although Wilson [Wil95] applied the update procedures as well as the GA originally in $[M]$ and only later [Wil98] was this switched to $[A]$, application in $[A]$ is now commonly used and gives better performance in all cases known to us. Thus we will use the $[A]$ notation in the following procedures although they are basically independent of the classifier set involved.

The reinforcement portion of the update procedure follows the pattern of Q-learning [SB98]. Classifier predictions are updated using the immediate reward and the discounted maximum payoff anticipated on the next time-step. The difference is that in XCS it is the prediction of a possibly general *rule* that is updated, whereas in Q-learning it is the prediction associated with an environmental *state-action* pair. Updates of classifier parameters other than prediction are unique to XCS. Note that in single-step problems, the prediction is updated using the immediate reward alone.

Each time a classifier enters the set $[A]$, its parameters are modified in the order: exp, p, ϵ, as, and F. Variations in the order are possible. The principle one is to exchange the p and ϵ updates. If prediction comes before error, the prediction of a classifier in its very first update immediately predicts the correct payoff and consequently the prediction error is set to zero. This can lead to faster learning in simpler problems but can be misleading in more complex ones. A more conservative strategy which puts the error update first, seems to work better on harder problems. The update of the action set size estimate is independent from the other updates and consequently can be executed at any point in time. While the updates of exp, p, ϵ, and as are straightforward, the update of F is more complex and requires more computational steps. Thus, we refer to another subprocedure. Finally, if the program is using action set subsumption, the procedure calls the *DO ACTION SET SUBSUMPTION* procedure. This procedure is very powerful and is able to eliminate a large number of classifiers in the action set in one step. Section 3.12 describes the procedure in detail.

Fitness Update The fitness of a classifier in XCS is based on the *accuracy* of its predictions. The *UPDATE FITNESS* procedure first calculates the classifier's

```
UPDATE SET([A], P, [P]):
 1 for each classifier cl in [A]
 2    cl.exp++
 3    //update prediction cl.p
 4    if(cl.exp < 1/β)
 5       cl.p ← cl.p + (P - cl.p) / cl.exp
 6    else
 7       cl.p ← cl.p + β * (P - cl.p)
 8    //update prediction error cl.ε
 9    if(cl.exp < 1/β)
10       cl.ε ← cl.ε + (|P - cl.p| - cl.ε) / cl.exp
11    else
12       cl.ε ← cl.ε + β * (|P - cl.p| - cl.ε)
13    //update action set size estimate cl.as
14    if(cl.exp < 1/β)
15       cl.as ← cl.as + (∑_{c∈[A]} c.n - cl.as) / cl.exp
16    else
17       cl.as ← cl.as + β * (∑_{c∈[A]} c.n - cl.as)
18 UPDATE FITNESS in set [A]
19 if(doActionSetSubsumption)
20    DO ACTION SET SUBSUMPTION in [A] updating [P]
```

accuracy κ using the classifier's prediction error ϵ. Then the classifier's fitness is updated using the *normalized accuracy* computed in lines 8-10.

```
UPDATE FITNESS([A]):
 1 accuracySum ← 0
 2 initialize accuracy vector κ
 3 for each classifier cl in [A]
 4    if(cl.ε < ε_0)
 5       κ(cl) ← 1
 6    else
 7       κ(cl) ← α * (cl.ε / ε_0)^{-ν}
 8    accuracySum ← accuracySum + κ(cl) * cl.n
 9 for each classifier cl in [A]
10    cl.F ← cl.F + β * (κ(cl) * cl.n / accuracySum - cl.F)
```

3.9 The Genetic Algorithm in XCS

The final sub-procedure in the main loop, *RUN GA*, is also the most complex. First of all, the action set is checked to see if the GA should be applied at all. In order to apply a GA the average time period since the last GA application in the set must be greater than the threshold θ_{GA}. Next, two classifiers (i.e. the parents) are selected by roulette wheel selection based on fitness and the offspring are created out of them. After that, the offspring are possibly crossed and mu-

tated. If the offspring are crossed, their prediction, error, and fitness values are set to the average of the parents' values. Finally, the offspring are inserted in the population, followed by corresponding deletions. However, if GA subsumption is being used, each offspring is first tested to see if it is subsumed by either of its parents; if so, that offspring is not inserted in the population, and the subsuming parent's numerosity is increased. (Besides checking if an offspring is subsumed by a parent, one could also check if it is subsumed by other classifiers in the action set, or even the population as a whole.)

RUN GA([A], σ, [P]):

```
 1 if(actual time t - Σ_{cl∈[A]} cl.ts * cl.n /  Σ_{cl∈[A]} cl.n > θ_GA)
 2    for each classifier cl in [A]
 3       cl.ts ← actual time t
 4    parent_1 ← SELECT OFFSPRING in [A]
 5    parent_2 ← SELECT OFFSPRING in [A]
 6    child_1 ← copy classifier parent_1
 7    child_2 ← copy classifier parent_2
 8    child_1.n = child_2.n ← 1
 9    child_1.exp = child_2.exp ← 0
10    if(RandomNumber[0,1) < χ)
11       APPLY CROSSOVER on child_1 and child_2
12       child_1.p ← (parent_1.p + parent_2.p) / 2
13       child_1.ε ← (parent_1.ε + parent_2.ε) / 2
14       child_1.F ← (parent_1.F + parent_2.F) / 2
15       child_2.p ← child_1.p
16       child_2.ε ← child_1.ε
17       child_2.F ← child_1.F
18    child_1.F ← child_1.F * 0.1
19    child_2.F ← child_2.F * 0.1
20    for both children child
21       APPLY MUTATION on child according to σ
22       if(doGASubsumption)
23          if(DOES SUBSUME parent_1, child)
24             parent_1.n++
25          else if(DOES SUBSUME parent_2, child)
26             parent_2.n++
27          else
28             INSERT child IN POPULATION [P]
29       else
30          INSERT child IN POPULATION [P]
31       DELETE FROM POPULATION [P]
```

While the sub-procedures *INSERT IN POPULATION* and *DELETE FROM POPULATION* are defined in section 3.10 and 3.11, respectively, the others are described in the following paragraphs.

Roulette-Wheel Selection The Roulette-Wheel Selection chooses a classifier for reproduction proportional to the fitness of the classifiers in set $[A]$. First, the sum of all the fitnesses in the set $[A]$ is computed. Next, the roulette-wheel is spun. Finally, the classifier is chosen according to the roulette-wheel result.

SELECT OFFSPRING($[A]$):

```
 1 fitnessSum ← 0
 2 for each classifier cl in [A]
 3   fitnessSum ← fitnessSum + cl.F
 4 choicePoint ← RandomNumber[0,1) * fitnessSum
 5 fitnessSum ← 0
 6 for each classifier cl in [A]
 7    fitnessSum ← fitnessSum + cl.F
 8    if(fitnessSum > choicePoint)
 9       return cl
```

Crossover The crossover procedure is similar to the standard crossover procedure in GAs. Implementations of XCS with one point and two point crossover were tested and resulted in approximately identical results. In the *APPLY CROSSOVER* procedure we show two-point crossover. The actions are not affected by crossover.

APPLY CROSSOVER(cl_1, cl_2):

```
 1 x ← RandomNumber[0,1) * (length of cl_1.C +1)
 2 y ← RandomNumber[0,1) * (length of cl_1.C +1)
 3 if(x > y)
 4    switch x and y
 5 i ← 0
 6 do{
 7    if(x ≤ i and i < y)
 8       switch cl_1.C[i] and cl_2.C[i]
 9    i++
10 }while(i < y)
```

A more sophisticated crossover could include assurance that x and y are different. However, even without such special checks the algorithm serves its purpose.

Mutation While crossover does not affect the action, mutation takes place in both the condition and the action. A mutation in the condition flips the attribute to one of the other possibilities. Mutation in the action changes it equiprobable to one of the other actions. Though more time-efficient methods are possible, we present here the simple one in which a die is flipped for each attribute. Since in XCS most of the time is spent in processes that operate on the whole population such as matching and deletion, the type of algorithm used for mutation

only slightly affects efficiency. Note that our mutation is somewhat restricted. Rather than allowing an attribute of the condition to change to any other attribute, we only allow changes either to # or to the specific value that matches the corresponding component of $\sigma(t)$. Thus the resulting condition still matches the current input. The effect (if the action is unchanged) is to search the current action set niche along the axis of specificity vs. generality. Less restricted mutation schemes are of course possible.

APPLY MUTATION(cl, σ):

```
 1 i ← 0
 2 do{
 3    if(RandomNumber[0,1) < μ)
 4       if(cl.C[i] = #)
 5          cl.C[i] ← σ[i]
 6       else
 7          cl.C[i] ← #
 8    i++
 9 }while(i < length of cl.C)
10 if(RandomNumber[0,1) < μ)
11    cl.A ← a randomly chosen other possible action
```

3.10 Insertion in the Population

The *INSERT IN POPULATION* procedure checks to see if the classifier to be inserted is identical in condition and action with a classifier already in the population. If so, the latter's numerosity is incremented; if not, the new classifier is added to the population.

INSERT IN POPULATION(cl, [P]):

```
1 for all c in [P]
2    if(c is equal to cl in condition and action)
3       c.n++
4       return
5 add cl to set [P]
```

3.11 Deletion from the Population

The deletion procedure realizes two ideas at the same time: (1) it assures an approximately equal number of classifiers in each action set, or environmental 'niche'; (2) it removes low-fitness individuals from the population. The following paragraphs describe the *DELETE FROM POPULATION* procedure and the *DELETION VOTE* sub-procedure.

Roulette-Wheel Deletion Like the selection procedure, the deletion procedure *DELETE FROM POPULATION* chooses individuals (for deletion) by roulette-wheel selection. But first, the procedure checks to see if the sum of classifier numerosities in $[P]$ is less than or equal to N. If so, the procedure exits, as deletion is unnecessary. Otherwise, the sum of all deletion votes is calculated, and a classifier is chosen for deletion. If the classifier is a macroclassifier and currently represents more than one classifier, then its numerosity is merely decreased by one. Otherwise, the classifier is completely removed from the population.

DELETE FROM POPULATION($[P]$):

```
if(∑_{c∈[P]} c.n ≤ N)
   return
avFitnessInPopulation ← ∑_{c∈[P]} c.F  / ∑_{c∈[P]} c.n
voteSum ← 0
for each classifier c in [P]
   voteSum ← voteSum + DELETION VOTE of c
                       with avFitnessInPopulation
choicePoint ← RandomNumber[0,1) * voteSum
voteSum ← 0
for each classifier c in [P]
   voteSum ← voteSum + DELETION VOTE of c
                       with avFitnessInPopulation
   if(voteSum > choicePoint)
      if(c.n > 1)
         c.n--
      else
         remove classifier c from set [P]
      return
```

The Deletion Vote As mentioned above, the deletion vote realizes niching as well as removal of the lowest fitness classifiers. The deletion vote of each classifier is based on the action set size estimate as. Moreover, if the classifier has sufficient experience and its fitness is significantly lower than the average fitness in the population, the vote is increased in inverse proportion to the fitness. In this calculation, since we are deleting one micro-classifier at a time, we need to use as fitness the (macro)classifier's fitness divided by its numerosity. The following *DELETION VOTE* procedure realizes all this.

DELETION VOTE(cl, $avFitnessInPopulation$):

```
vote ← cl.as * cl.n
if(cl.exp > θ_del and cl.F / cl.n  < δ * avFitnessInPopulation)
   vote ← vote * avFitnessInPopulation / ( cl.F / cl.n)
return vote
```

3.12 Subsumption

Two subsumption procedures were introduced into XCS in [Wil98]. The first, 'GA subsumption', checks an offspring classifier to see if its condition is logically subsumed by the condition of an accurate and sufficiently experienced parent. If so, the offspring is not added to the population, but the parent's numerosity is incremented. The idea is that such an offspring cannot improve the system's performance, since everything it accomplishes is accomplished just as well by the subsuming parent. GA subsumption, if enabled, occurs within the procedure *RUN GA*. It is detailed within that procedure, and is not called as a sub-procedure (though it could be). However, the sub-procedure *DOES SUBSUME* is called, and is described in this section.

The second subsumption procedure, 'action set subsumption', if enabled, takes place in every action set. It has a purpose similar to GA subsumption but is different and independent of it. The action set is searched for the most general classifier that is both accurate and sufficiently experienced. Then all other classifiers in the set are tested against the general one to see if it subsumes them. Any classifiers that are subsumed are eliminated from the population.

DO ACTION SET SUBSUMPTION($[A]$*,* $[P]$*):*

```
 1 initialize cl
 2 for each classifier c in [A]
 3    if(c COULD SUBSUME)
 4       if(cl empty or number of # in c.C > number of # in cl.C
           or (number of # in c.C = number of # in cl.C and
                                        RandomNumber[0,1) < 0.5))
 5          cl ← c
 6 if(cl is not empty)
 7    for each classifier c in [A]
 8       if(cl IS MORE GENERAL than c)
 9          cl.n ← cl.n + c.n
10          remove classifier c from set [A]
11          remove classifier c from set [P]
```

Subsumption of a classifier For a classifier to subsume another classifier, it must first be sufficiently accurate and sufficiently experienced. This is tested by the *COULD SUBSUME* function. Then, if a classifier could be a subsumer, it must be tested to see if it has the same action and is really more general than the classifier that is to be subsumed. This is the case if the set of situations matched by the condition of the potentially subsumed classifier form a proper subset of the situations matched by the potential subsumer. The *IS MORE GENERAL* procedure accomplishes this. The *DOES SUBSUME* procedure combines all the requirements.

COULD SUBSUME(cl):

```
1 if(cl.exp > θ_sub)
2    if(cl.ε < ε_0)
3       return true
4 return false
```

IS MORE GENERAL(cl_{gen}, cl_{spec}):

```
1 if (the number of # in cl_gen.C ≤ the number of # in cl_spec.C)
2    return false
3 i ← 0
4 do{
5    if(cl_gen.C[i] ≠ # and cl_gen.C[i] ≠ cl_spec.C[i])
6       return false
7    i++
8 }while(i < length of cl_gen.C)
9 return true
```

DOES SUBSUME(cl_{sub}, cl_{tos}):

```
1 if(cl_sub.A = cl_tos.A)
2    if(cl_sub COULD SUBSUME)
3       if(cl_sub IS MORE GENERAL than cl_tos)
4          return true
5 return false
```

4 Summary

This paper has revealed the processes inside XCS as well as the problem interaction. We hope that presentation of the algorithm in pseudo code with explanations will lead to deeper understanding of XCS and simplify research. Moreover, the modular structure should enable the reader to program the system in any programming language quite easily. However, our description only includes the basic framework of XCS. Starting from this baseline, we encourage modification as well as enhancement of the system, and welcome feedback on progress and results.

Acknowledgments

The work was sponsored by the Air Force Office of Scientific Research, Air Force Materiel Command, USAF, under grants F49620-97-1-0050 and F49620-00-0163. Research funding for this work was also provided by the National Science Foundation under grant DMI-9908252. Support was also provided by a grant from the U. S. Army Research Laboratory under the Federated Laboratory Program, Cooperative Agreement DAAL01-96-2-0003. The U. S. Government is authorized to reproduce and distribute reprints for Government purposes notwithstanding any copyright notation thereon. The authors acknowledge support from the Au-

tomated Learning Group at the National Center for Supercomputer Applications in Urbana-Champaign.

The views and conclusions contained herein are those of the authors and should not be interpreted as necessarily representing the official policies or endorsements, either expressed or implied, of the Air Force Office of Scientific Research, the National Science Foundation, the U. S. Army, or the U. S. Government.

References

[BW01] Martin V. Butz and Stewart W. Wilson. An algorithmic description of XCS. In P. L. Lanzi, W. Stolzmann, and S. W. Wilson, editors, *Advances in Learning Classifier Systems, LNAI 1996*, Berlin Heidelberg, 2001. Springer-Verlag.

[DC98] Marco Dorigo and Marco Colombetti. *Robot shaping: An experiment in behavior engineering*. Intelligent Robotics and Autonomous Agents. MIT Press, Cambridge, MA, 1998.

[Kov97] Tim Kovacs. XCS classifier system reliably evolves accurate, complete, and minimal representations for boolean functions. In Roy, Chawdhry, and Pant, editors, *Soft Computing in Engineering Design and Manufacturing*, pages 59–68. Springer-Verlag, 1997.

[Kov99] Tim Kovacs. Deletion schemes for classifier systems. In Wolfgang Banzhaf, Jason Daida, Agoston E. Eiben, Max H. Garzon, Vasant Honavar, Mark Jakiela, and Robert E. Smith, editors, *Proceedings of the Genetic and Evolutionary Computation Conference (GECCO-99)*, pages 329–336, San Francisco, CA, 1999. Morgan Kaufmann.

[Lan97] Pier Luca Lanzi. A study of the generalization capabilities of XCS. In T. Baeck, editor, *Proceedings of the Seventh International Conference on Genetic Algorithm*, pages 418–425, San Francisco, California, 1997. Morgan Kaufmann.

[Lan99a] Pier Luca Lanzi. An analysis of generalization in the XCS classifier system. *Evolutionary Computation*, 7(2):125–149, 1999.

[Lan99b] Pier Luca Lanzi. Extending the representation of classifier conditions. Part I: From binary to messy coding. In Wolfgang Banzhaf, editor, *Proceedings of the Genetic and Evolutionary Computation Conference (GECCO-99)*, pages 337–344, San Francisco, CA, 1999. Morgan Kaufmann.

[Lan99c] Pier Luca Lanzi. Extending the representation of classifier conditions. Part II: From messy coding to S-expressions. In Wolfgang Banzhaf, editor, *Proceedings of the Genetic and Evolutionary Computation Conference (GECCO-99)*, pages 345–352, San Francisco, CA, 1999. Morgan Kaufmann.

[Lan99d] Pier Luca Lanzi. An extension to the XCS classifier system for stochastic environments. In Wolfgang Banzhaf, editor, *Proceedings of the Genetic and Evolutionary Computation Conference (GECCO-99)*, pages 353–360, San Francisco, CA, 1999. Morgan Kaufmann.

[LSW00] P. L. Lanzi, W. Stolzmann, and S. W. Wilson, editors. *Learning classifier systems: From foundations to applications*. LNAI 1813. Springer-Verlag, Berlin Heidelberg, 2000.

[LW00] Pier Luca Lanzi and Stewart W. Wilson. Toward optimal classifier system performance in non-markov environments. *Evoultionary Computation*, 8(4):393–418, 2000.

[SB98] Richard S. Sutton and Andrew G. Barto. *Reinforcement learning: An introduction*. MIT Press, Cambridge, MA, 1998.

[Wil95] Stewart W. Wilson. Classifier fitness based on accuracy. *Evolutionary Computation*, 3(2):149–175, 1995.

[Wil98] S. W. Wilson. Generalization in the XCS classifier system. In John R. Koza, Wolfgang Banzhaf, Kumar Chellapilla, Kalyanmoy Deb, Marco Dorigo, David B. Fogel, Max H. Garzon, David E. Goldberg, Hitoshi Iba, and Rick Riolo, editors, *Genetic Programming 1998: Proceedings of the Third Annual Conference*, pages 665–674, San Francisco, 1998. Morgan Kaufmann.

[Wil00] Stewart W. Wilson. Get real! XCS with continuous-valued inputs. In P. L. Lanzi, W. Stolzmann, and S. W. Wilson, editors, *Learning Classifier Systems: From Foundations to Applications, LNAI 1813*, pages 209–220, Berlin Heidelberg, 2000. Springer-Verlag.