A Search Based Approach for Overlapping Concept Boundaries

Nicolas Gold Mark Harman Zheng Li
Kiarash Mahdavi
King's College London
Computer Science Department
Strand, London WC2R 2LS, England, United Kingdom

Abstract

This paper presents techniques to integrate boundary overlap into concept assignment using Plausible Reasoning. Heuristic search techniques such as Hill climbing and Genetic Algorithms are investigated. A new fitness measure appropriate for overlapping concept assignment is introduced. The new algorithms are compared to randomly generated results and the Genetic Algorithm is shown to be the best of the proposed search algorithms in terms of the quality of concept binding, as measured by the fitness function. The fixed boundary Hypothesis Based Concept Assignment technique is compared to the new algorithms. The Genetic Algorithm and Hill climbing are found to consistently produce stronger concepts than Hypothesis Based Concept Assignment.

1. Introduction

Program comprehension is one of the most expensive activities in software maintenance and many tools and techniques have been created to reduce the time and expense involved. Concept assignment techniques, such as Hypothesis Based Concept Assignment (HBCA) have been successfully employed to assign descriptive terms to source code as a means for program comprehension [3, 6]. The resulting concepts created by concept assignment techniques such as HBCA are distinct, non overlapping segments of code, which relate to computational intent. However, complete distinct localisation of concept within code is perhaps too rigid an assumption within real programs [17]. Concepts created without this assumption may be a better representation of computational intent in the code. Figure 1 contains an example of a concept overlap in a code fragment of COBOL II [4]. In this example the first, third and fifth lines indicate a **Print** concept and the second, fourth and sixth lines indicate a **Write** concept. It is impossible to determine where one concept ends and the other starts. It could also be reasoned that the last two lines indicate a **Call** concept.

```
MOVE  'EXAMPLE' TO PRINT-LL.
MOVE  POLICY-NUM TO OUT-PNUM.
MOVE  '13' TO PRINT-CC
MOVE  SCHEME-REF TO OUT-SERF.
CALL  'PRINT'  USING P-PRINTLINE
CALL  'WRITE' USING OUT-REC
```

Figure 1. Example of overlapping concepts

HBCA is a plausible reasoning technique with a linear growth in computation cost, which merits its application for large program studies [1]. This made it a suitable candidate for studying a large number of programs of varying size within this experiment. In addition, the concept binding mechanism (explained further in section 2) could be easily adapted as a means to drive the concept binding process for the new search based techniques.

This paper contains the details of an investigation into using search based approaches for concept assignment that allow overlap of concept bindings. Genetic Algorithms (GA) and Hill Climbing (HC) search algorithms were used. GA's ability to tackle complex and fuzzy problems made it a suitable candidate. HC was selected to determine if solutions achieved by using a less complex and quicker local search were able to compete with the more computationally intensive results from the GA. The results, were compared and analysed. They demonstrate the improvements gained due to the use of overlapping concept assignment with the GA and HC over the original HBCA algorithm. The overall contributions of this paper can be summarised as:

- Formulating the overlapping concept boundaries as a search problem.
- Devising search algorithms to allow overlapping concept boundaries for concept assignment.
- Devising a Fitness Function suitable for evaluating

concepts with overlapping boundaries

- Devising a solution structure to allow search for concepts with overlapping boundaries.
- Empirical study of the algorithms and analysis of the results.

2. Hypothesis-Based Concept Assignment

This section contains a brief explanation of the HBCA algorithm, since HBCA algorithms and its results are drawn upon to explore the effect of allowing overlapping boundaries in concept assignment. HBCA requires a library or knowledge base. This library is a semantic network, which is composed of *Concepts* and *Indicators*. Indicators are evidence for concepts within the implementation. Concepts are the terms nominated by the user to describe items or activities in the domain. The library also includes relationships between concepts, which are used to identify composite or specialised concept binding. A concept may take the form of an Action or Object. Action concepts carry out operations ,for example **Write** is an Action concept. Object concepts are concepts which can be acted upon by Action concepts and their presence together may suggest the existence of a composite concept, for instance the Object concept **File** and **Write** can create the composite **Write File** concept. The more general forms of object concepts are regarded as primary while the more specialised form are regarded as secondary. Composite concept may be created by using identified relationships within the library between Action and Object/Spesialised concepts. The HBCA algorithm is summarised in three stages. *Hypothesis Generation*, *Segmentation* and *Concept Binding*.

2.1. Hypothesis Generation

Hypothesis Generation draws on source code as input. The library is utilised to scan the source code for indicators of various concepts. For each matching concept, a hypothesis is generated and stored. The list of hypothesis is ordered according to the position of the indicators in the source code. Where all the further stages of Segmentation and Concept Binding are carried out on this created *Hypothesis List*. Figure 2 contains an example of a generated *hypothesis list*. The created Hypothesis List is also the input for the search based algorithms.

2.2. Segmentation

Segmentation stage attempts to create distinct, disjoint segments within a Hard Segment. Hard Segments are natural segment boundaries such as procedure divisions. A Self Orgainsing Map (SOM) creates segments of high conceptual focus according to the distribution of the the Action concepts within a Hard Segment.

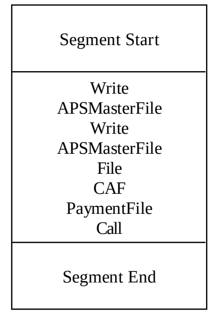

Figure 2. Example of a generated Hypothesis List

2.3. Concept Binding

Concept Binding is carried out by the Concept Assigner. The Concept Assigner evaluates each segment in term of concept occurrence according to simple (Action concepts) and composite (Action/Object/specialised object) concepts. The library is used to determine the possible composite concepts within each segment. The strength of evidence for a concept is equivalent to the number of hypothesis that could indicate the presence of that concept. The Concept Assigner also requires the presence of at least one action concept in addition to a user defined minimum number of hypothesis (minimum evidence) to create a *Concept Binding*. Assuming these conditions are satisfied, the concept with the strongest evidence is selected as the winner. A set of disambiguation rules are applied to select a winner in case of ties. An in depth analysis of these rules can be located in Gold's PhD thesis [6]. The segments are then bound to the winning concepts and highlighted in code.

3. Defining the Search Problem

A Search Problem is the algorithmic identification of of a solution from a solution space. As discussed in Subsection 2.1, the input for this search algorithms is the Hypothesis List generated by the HBCA algorithm. The problem therefore can be defined as searching for Segments of Hypothesis in each Hypothesis List according to predetermined fitness criteria such that each Segment has the following attributes:

- Each Segment contains 1 or more neighbouring Hypothesis.
- There are no duplicate Segments.

The search fitness criteria's aims are twofold:

- Guide the search to finding Segments of strongest evidence.
- Binding as many of the Hypothesis within the Hypothesis List without compromising the Segments strength of evidence.

What follows in this Section is an investigation of the implications of the above definitions on the size of the search in Subsection 3.1. In addition, a Fitness Function based on the above guidelines is proposed in Subsection 3.2. An overview of GA and HC algorithms and their more specific design with respect to this experiment are explained in Subsections 3.3 and 3.4. Finally a solution structure for the algorithms is presented in Subsection 3.5. This Subsection also explores some of the design and implementation issues raised and dealt with as a result of utilising the proposed solution structure.

3.1. Size of the Search Landscape

In this section the size of the search (number of possible solutions) and its growth according to the number of Hypothesis within within the Hypothesis List are analysed. The number of possible Segments which can be created within a Hypothesis List ,according to the definition in Section 3, is given explicitly by the following:

$$s = \frac{h(h+1)}{2}$$

Where s is the number of Segments and h represents the number of Hypothesis. The size of the search (the number of possible solutions) is the number of possible combinations of these Segments:

$$c = \sum_{k=1}^{s} \binom{s}{k}$$

c is the number of possible Segment combinations. $\binom{n}{k}$ represents binomial coefficient, which is the number of possible subsets of size k within the set of segments s. c is calculated explicitly by the powerset of s minus the empty set:

$$c = 2^s - 1$$

Figure 3 demonstrate the rapid resulting increase in the search size c from steady increase in the number of Hypothesis h. This rapid increase makes the use of a population based search algorithm like a GA all the more appealing due to their previous success in tackling large search spaces examples of which can be found in [13, 21, 22].

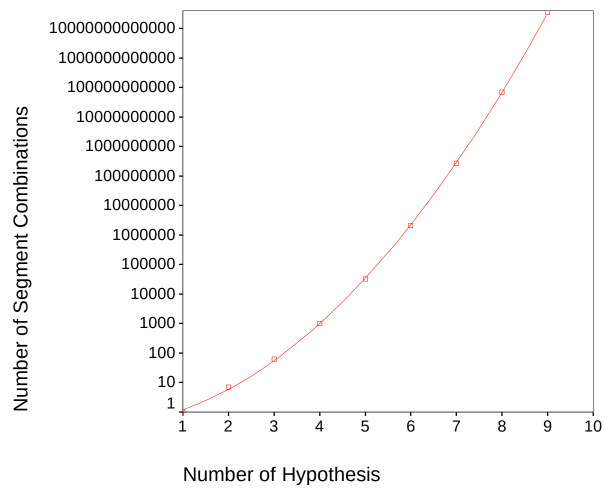

Figure 3. Search size increase against increasing number of Hypothesis in the Hypothesis List.

3.2. Fitness Function

For the Fitness Function to effectively guide the search, it must be able to evaluate each solution according to the strength of evidence and Hypothesis List coverage. The first step involves the recognition of the strongest concept within each Segment of a particular solution. This is achieved by following the same process as the HBCA's *Concept Binding* previously discussed in Subsection 2.3.

The overall fitness is then evaluated in order to find the Segmentation Strength in addition to the Hypothesis List coverage. The Segmentation Strength is the combination of Inner fitness and the Potential Fitness of each Segment. The inner fitness of each Segment is assessed as:

$$fit_i = signal_i - noise_i$$

Where fit_i signifies the Inner Segment Fitness, $signal_i$ represents the Signal level or number of Hypothesis within the Segment that contribute to the winner and $noise_i$ represents the Noise level or the number of Hypothesis within the Segment that do not contribute to the winner. The Inner Segment Fitness results in recognition of higher Fitness for Segments with more evidence indicating their winning concept. In addition Each Segment is evaluated with respect to the entire Segment Hypothesis List:

$$fit_p = signal_i - signal_p$$

The Potential Segment Fitness, fit_p is evaluated by taking account of $signal_p$, the number of Hypothesis outside of the Segment that have could contribute to the Segment's winning concept if they were included in the Segment. This facet of the Segment Fitness effectively guides the search to creating Segments of larger size in order to incorporate as much of the Signal as possible. The Overall Segment Fitness is evaluated by combining these Inner and Potential Segment Fitness into an Overall Segment Fitness:

$$segfit = fit_i + fit_p$$

The Overall Segment Fitness ($segfit$) attempts to guide the search at a local Segment level to larger Segments according to higher Signal Potential while preserving a low Noise level. Finally the Total Segment Fitness is calculated as:

$$totsegfit = \sum_{s=1}^{n} = segfit(s)$$

where s is represents Segments within the solution.

Hypothesis List Coverage is the second facet of the Fitness calculation. Increased coverage of the Hypothesis List results in further coverage of the original program code, which could potentially improve program comprehension. Hypothesis List Coverage is defined as:

$$hc = h - hn$$

where h is the number of Hypothesis within the Hypothesis List and hn the number of Hypothesis not covered by any Segments within the solution. For a just comparison of solutions, given that solutions may have a varying number of Segments and Coverage, a normalised version of the Fitness is evaluated:

$$\frac{totsegfit + hc}{2totseglength + h}$$

where

$$totseglength = \sum_{s=1}^{n} seglength(s)$$

and $seglength(s)$ is the number of Hypothesis in Segment s.

3.3. Genetic Algorithms

GAs are a collection of heuristic population based evolutionary search techniques. Traditionally individual solutions with a GA population are also referred to as Chromosomes and their constituent bits as Genes [11]. The Genes represent a coding of the search parameters rather than directly used parameters. GA search starts with a random population of potential solutions. It then employs evolutionary aspired operators whilst guided by a Fitness Function to evolve fitter solutions. Amongst the many subtly varied definitions of a GA the following mechanisms are agreed upon [8, 15, 19]:

Selection is the probability based sampling of the current population , which is guided by a set of fitness objectives (fitness function). The selected solutions participate in creating the next generation of solutions using GA search operators.

Crossover is the primary search operator. It involves the recombination of pairs of solutions picked during Selection as parents in order to create offsprings for the next generation of solutions.

Mutation is the secondary search operator. It involves the random selection and change of Genes in the a newly created population with the aim of reducing population stagnation.

Selection, Crossover and Mutation are used to create subsequent generations of populations until the Stopping Condition is met where the algorithm terminates. The Stopping Condition maybe based heuristically according, where after a predetermined number of generations the search terminates or based deterministically.

Population size, Selection, Crossover and Mutation are governed by a set of heuristics. Crossover Rate determines the probability of recombination for a pair of selected individuals. Selected individuals which are not recombined due to Crossover Rate are copied directly into the new population. Tournament Selection is explained in this section as it is used as the Crossover operator for the proposed GA. Initially a random pair of Chromosomes are selected for Tournament Selection. The Chromosome with the higher fitness may then be selected to participate in Crossover and Mutation according to the Tournament Coefficient heuristic. The Tournament Coefficient is the probability of the fitter Chromosome being selected and is usually strongly set in favour of the fitter Chromosome. Mutation Rate determines the probability of a Mutation per Gene(bit) of a solution of the new population. This probability is normally set to be very small to avoid the search from deteriorating into a random search. The Crossover and Mutation mechanisms are also dependent on the defined structure of the chromosome, the proposed structure in addition to its implications for these operations are discussed further in Subsection 3.5.

3.4. Hill Climbing Algorithms

Hill Climbing is a local search technique. It stars from a single randomly created solution. A predefined set of Mutations are then used to define a set of potentially fitter solutions. The set of solutions created by the Mutation operator are referred to as the neighbouring solutions. The search at each stage selects a fitter neighbouring solution and examines further neighbouring solutions from the newly discovered solution. The number of considered neighbouring solutions is a heuristic that determines the minimum number of examined before a selection decision is made. Due to the local nature of the search, it is possible for the starting solution to quickly find local Optimal solution of relatively low fitness. The search attempts to reduce this effect by restarting from a new random solution or by using the characteristics of the current solution to create a new solution. The search ends when no fitter neighbours can be found and the Sopping Condition has been reached. The Stopping Condition search may be determined heuristically by limiting the number of allowed restarts or algorithmically when no fitter solution for a restart can be found by using the restart mechanism. The proposed HC for this experiment uses characteristics of the final solution for the restart operation and stop when it can not create better solutions by using this restart method. The restart operation and permitted Mutations are dependent on the defined structure of the solution. The proposed solution structure and the exact nature of the restart and Mutations are further discussed in the *Solution Structure* Subsection (Subsection 3.5).

3.5. Solution Structure

The scope of resulting solutions are affected by what constitutes a solution in the HC and chromosomes in the GA population. In this case, each solution needs to represent all the discovered Segments and be flexible enough to allow independent Segment boundaries and a variable number of Segments. The proposed solution defines a Segment as a pair of values where each values represent the location of start and end Hypothesis within the ordered Hypothesis List. Since the number of Segments within a Hypothesis List in not predetermined, the length of a potential solution can also vary. A messy GA chromosome structure was chosen as a suitable representation [9] since it allows the Chromosome to have variable length. In the proposed solution a Chromosomes is made up of a set of one or more Segments representations. For future ease of reference, the Segment representations will be referred to as Segment Pairs.

A difficulty that was detected during implementation of this representation is the potential for an unmanageable increase in the size of solutions. Although the problem definition as explained in Section 3 results in the elimination of all duplicates Segment Pairs, this never the less leaves a potentially large number of Segment Pairs. The Fitness Function was exploited in this case to reduce this size even further. In the proposed solution all Segments with the same winning concept that overlap are compared and all but the fittest Segment are removed from the solution.

Crossover and Mutation operations are also directly dependent on the solution structure. The early GA experimentations with the Crossover involved random selection and recombination of parent's Segment Pairs to create new Segment Pairs in the offspring. This resulted in rapid deterioration of the search into almost randomly generated new solutions, where a large amount of useful information was destroyed in the recombination process. The proposed implemented solution utilises the location of the segment pairs, where only segment pairs of overlapping locations are recombined and the remaining are copied to the new chromosome. Figure 4 contains an example of the recombination process on two chromosomes and the resulting offspring.

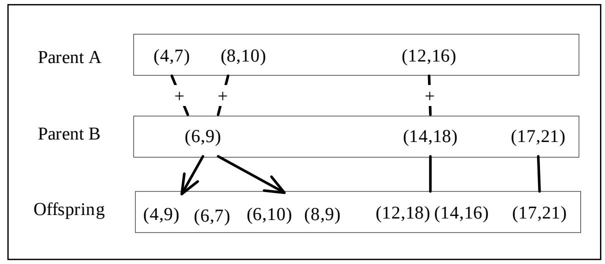

Figure 4. Example Crossover operation on GA chromosomes

The GA Mutation operation on the proposed solution structure is different from HC Mutation. The GA Mutation operator is a secondary search operator, which is primarily concerned with population stagnation. Therefore the Mutation operator can randomly replace any Hypothesis location within any Segment Pair with any other valid Hypothesis location with the concern for causing the search to become overly randomised. As a result this Mutation model is used for the GA, where a Mutation can occur on each Hypothesis location according to the Mutation Rate, where a Mutation results in the replacement of a Hypothesis location value in a Segment Pair with a random and valid Hypothesis location value. Conversely such Mutations would cause a local search technique such as Hill Climbing to become akin to a random search. To reduce this effect, the proposed HC Mutation operator generates new solutions by selecting a Segment Pair and increasing or decreasing one of the location values by a single increment. The resulting Mutations are a set of similar neighbouring solutions which can be used to describe the local search landscape of the Hill Climbing algorithm. Finally the proposed HC takes advantage of the proposed GA Crossover operation GA for the Restart mechanism as discussed in Subsection 3.4. This entails recombination of all Segment Pairs in order to create new Segments Pairs, which are then added to the current Solution if their inclusion results in an improvement to the Fitness value.

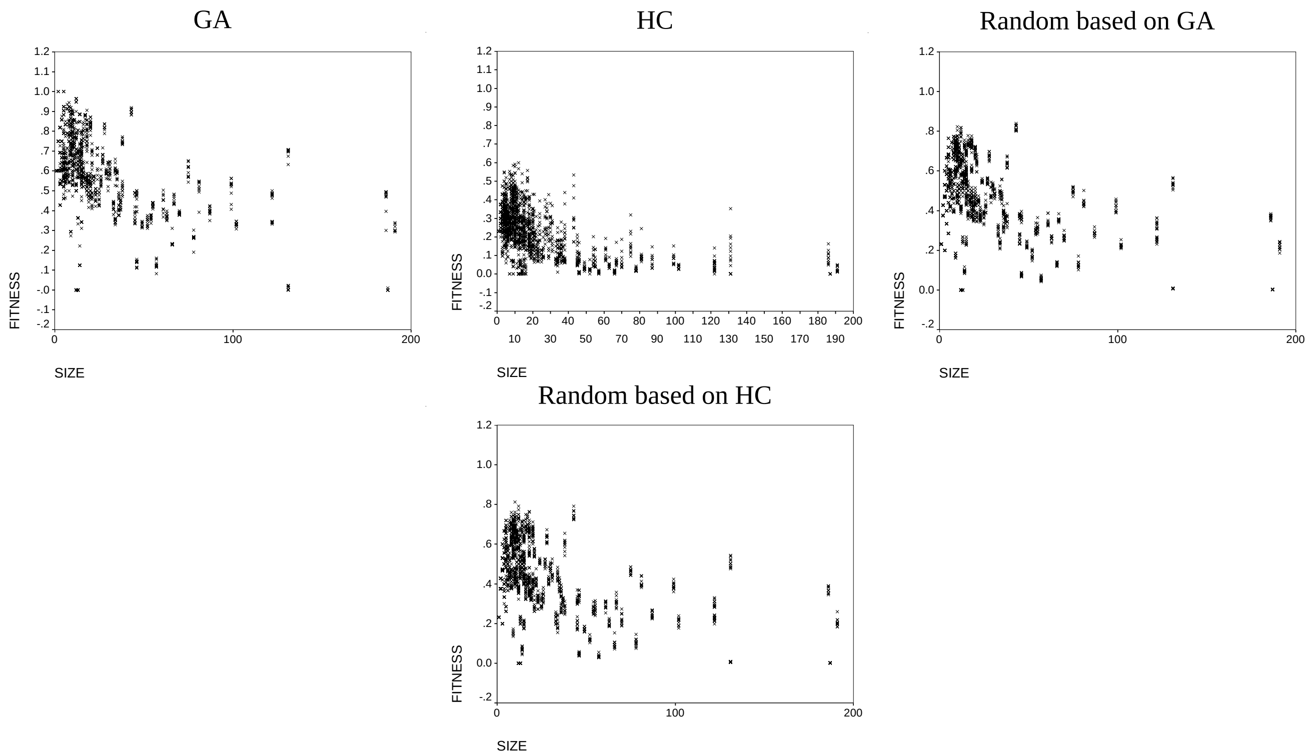

Figure 5. Scatter graphs of fitness against hard segment size.

4. Empirical Study

An empirical study was carried out to identify the best of the proposed algorithms for Concept Assignment that allow overlapping Concept Boundaries based on the proposed Fitness discussed in Subsection 3.2. 21 COBOL II programs were included in the study. A set of Hypothesis lists were generated by using the HBCA Hypothesis Generator discussed in Subsection 2.1. The generated Hypothesis Lists have a size ranging from 1 to 191 Hypothesis. Due to the probabilities involved in the generation of the initial population for the GA and initial solution for the HC and the probabilities involved in the search operators, it is possible to generate results of uncharacteristic low fitness for these algorithms. In order to to better evaluate characteristic solutions, 10 GA and HC runs were carried out per Hypothesis List.

The set of heuristic values for the GA and HC were derived by trials and experimentation on some of the smaller Hypothesis List. The GA population consist of 100 Chromosomes, which are created randomly for the initial population. This involves the creation of a random number of Genes or Segment Pairs. The number of Genes is set between a minimum of 5 to a maximum based on the number of Hypothesis in the Hypothesis List. The Tournament Selection's Coefficient were set to 0.99 and Crossover and Mutation Rates were set to 0.8 and 0.001 respectively. The GA search terminates when the average fitness of the population does not 50 generations. The HC algorithm initial solutions are produced by following the GA initial Chromosome creation mechanism. The number of considered neighbours are set to 1 and the stopping condition is met when no better solution can be achieved by using the restart mechanism explained in Subsection 3.5.

The GA and HC results were also compared to sets of randomly generated solutions for each Hypothesis List. These solutions were created according to the solution structure described in Subsection 3.5. Number of generated Random solutions for each Hypothesis List was determined by the number of evaluated solutions by the GA and HC algorithms for each Hypothesis List. Comparison of the GA, HC and Random results are discussed in Subsection 4.1.

HBCA was also used for Concept Assignment on the 21 COBOL II programs. The minimum evidence level for Concept Binding was set at 3 Hypothesis (details of HBCA Concept Binding and evidence level have been discussed in Subsection 2.3). Analysis was carried out on the results from the proposed GA and HC algorithms and the HBCA results to find the best algorithm. Due to the different HBCA search criteria the comparison is based on a different measure. This along with the results of the study are discussed in Subsection 4.2.

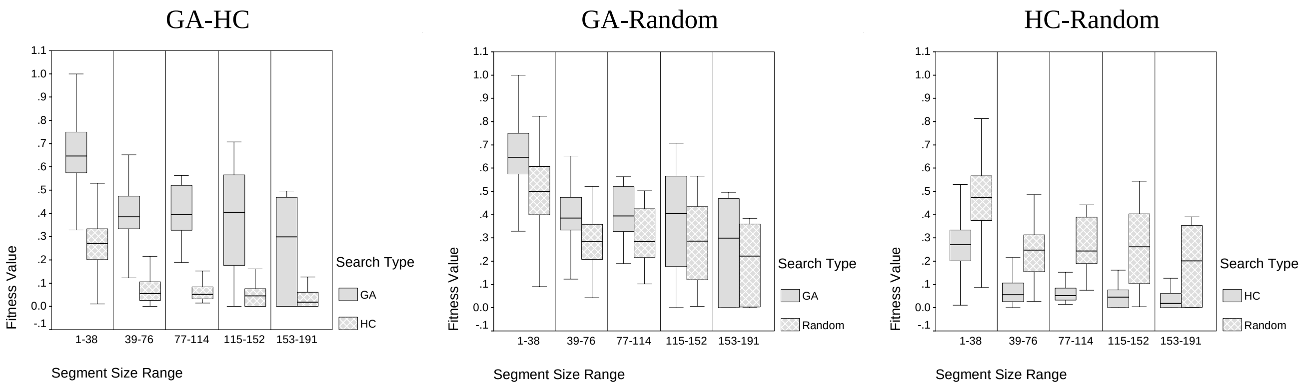

Figure 6. Boxplots of GA, HC and Random Fitness Results compared across increasing Hypothesis List size range.

The results are presented as scatter graphs and boxplots. The scatter graph's vertical axis represents the fitness value and its horizontal axis represents increasing Hypothesis List size. Similarly the boxplots vertical axis represents fitness values, however to reduce clutter in the presentation of boxplots, caused by the large variety in Hypothesis List sizes, it was necessary for the boxplots to be drawn against increasing ranges of Hypothesis List size. The 5 increasing ranges used are 1 to 38, 39 to 76, 77 to 114, 115 to 152 and 153 to 191. For example the range 1 to 38 represents all the fitness values resulting from Hypothesis List sizes of between 1 to 38 (inclusive) Hypothesis. Each boxplot represents the distribution of Fitness values for a particular Hypothesis List size range. The length of the box corresponds to the interquartile range and contains 50% of cases. The line across the inside of the box indicates the median value. The protruding lines (whiskers) represent the smallest and largest fitness values that are not outliers, where outliers are values which are 1.5 box-lengths away from the edge of the box.

4.1. GA, HC and Random Results

The results from the GA, HC and Random experiment were compared based on their fitness values, which were calculated by the Fitness Function described in Subsection 3.2. Scatter graphs labeled GA, HC, Random based on GA and Random based on HC in Figure 5 contain the Fitness distribution ordered by Hypothesis List size for GA, HC and GA and HC random search respectively. The distribution of GA Fitness results according to Figure 5 is similar to the other distributions in that Figure but is present at a higher Fitness level. On the other hand, the HC results are clearly of inferior Fitness to all other results presented by the scatter graphs in Figure 5. The boxplot of Fitness values for paired algorithms against increasing Hypothesis List size range are overlayed in Figure 6. The superior GA results compared to HC and Random is also conveyed by this Figure. The GA-HC and HC-Random graphs in this Figure once again highlight the inferior HC Fitness results compared to the other results. Not surprisingly the results also demonstrate the increasing difficulty for all search algorithms as the Hypothesis List size increases. This observation corresponds to the rapid increase in the search space discussed in Subsection 3.1. The HC also has the smallest distribution of results compared to the GA and Random results. The implications of this observation are further discussed in Section 6.

Pair-wise comparison of the GA, HC and Random by the *Wilcoxen Signed Rank Test* was used to ascertain the statistical significance in the observed strength of GA and the weakness of the HC results in terms of Fitness values. The *Wilcoxen Signed Rank Test* reported the level of significance to be less that 0.0005 when comparing GA results against HC and GA results against Randomly created results. This level of significance represents a statistically significant improvement to the Fitness for the GA compared to the HC and Randomly generated solutions. The *Wilcoxen Signed Rank Test* reported a level of significance of below 0.0005 meaning a statistically signinficant worsening of HC results compared to the Randomly generated solutions.

4.2. GA, HC and HBCA Results

As described in Section 2, The HBCA algorithm carries out Segmentation based on a different set of criteria to GA and HC algorithms. For a more impartial comparison between these algorithms, the results are evaluated based on the Signal to Size ratio, where the Signal represents the number of Hypothesis within a Segment that contribute to the winning concept and Size represents the number of Hypothesis within that Segment.

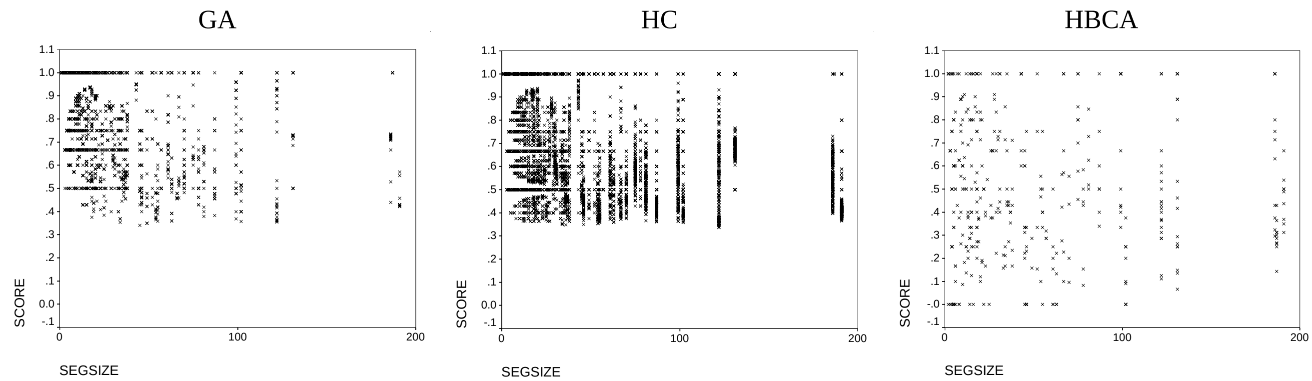

Figure 7. Scatter graphs of Signal to Size ratio against hard segment size.

Scatter graphs in Figure 7 displays the distribution of Signal to Size ratios of created Segments across increasing Hypothesis List size for the GA, HC and HBCA algorithms. Most noticeable from these graphs is the lack of solutions with low Signal to Noise ratios for the scatter graphs of GA and HC when compared to HBCA. The reasons for this observation are difficult to determine due to the different fitness criteria that the GA and HC algorithms use compared to the HBCA algorithm. The GA-HBCA boxplots in Figure 8 further illustrate characteristically better Signal to Noise ratios achieved by the GA algorithm compared to HBCA across the range of of Hypothesis Lists. The HC results in the HC-HBCA boxplots in Figure 8, although not as clearly improved as the GA results, are generally better when compared to HBCA. The graphs in Figure 8 also display consistently higher Signal to Noise ratio of GA results in comparison to HC.

Pair-wise comparison of the GA, HC and HBCA was carried out by using the *Wilcoxen Signed Rank Test* to determine if the strength of GA and HC results against HBCA were significant in terms of Signal to Size ratio. The test reported a significant difference of less than 0.0005 for all of these comparisons. This implies the GA and HC Signal to Size ratios were significantly better than HBCA. Further *Wilcoxen Signed Rank Test* between GA and HC also yielded a significance difference of below 0.0005, meaning in terms of Signal to Noise ratio, the GA results were also significantly better than the HC.

5. Related Work

Concept assignment has been defined as "...a process of recognising concepts within a computer program and building an 'understanding' of the program by relating recognised concepts to portions of the program, its operational context and to one other [1]."

The two major research issues of concept assignment have been identifies by Tilly et al. as *Segmentation* and *Concept Binding* [23]. Segmentation is the process of grouping pieces of conceptual information generated from the source code. Concept binding involves the analysis of segments in order to determine the most plausible concept assignment for each segment [5]. The segmentation and concept binding process are intricately and naturally linked. The location and size of the segment determines the assigned concept. The strength of the assigned concept on the other hand determines the quality of segmentation. Another problem directly involved in segmentation of software is the presence of overlapping concepts within the software.

Concept assignment techniques are carried out by intellegent agent tools. They traditionally fall within the following categories [1].

1. Domain specific, rule based, model driven systems that answer specific questions. These depend on manually created databases which describe the software system. This approach is exemplified by the LaSSIE System [2].

2. Plan driven, algorithmic and based on a precise set of understanding and recognition rules. Examples of this method are Programmer's Apprentice [20] and GRASPR [25].

3. Model driven systems which use plausible reasoning. Examples of this technique are DM-TAO [1], IRENE [14] and HBCA [3, 6].

The tools which employ approaches in first and second category are capable of completely deriving concept within small scale programs but due to their overwhelming computational cost are not suitable for larger-scale programs [1]. Conversely, the third approach have a linear computational growth for increased program size, but suffer from imprecision in results [1].

HBCA [3, 6] plausible reasoning techniques has recently been proposed as a means for more complex reveres engineering and software testing. This involves the use of pro-

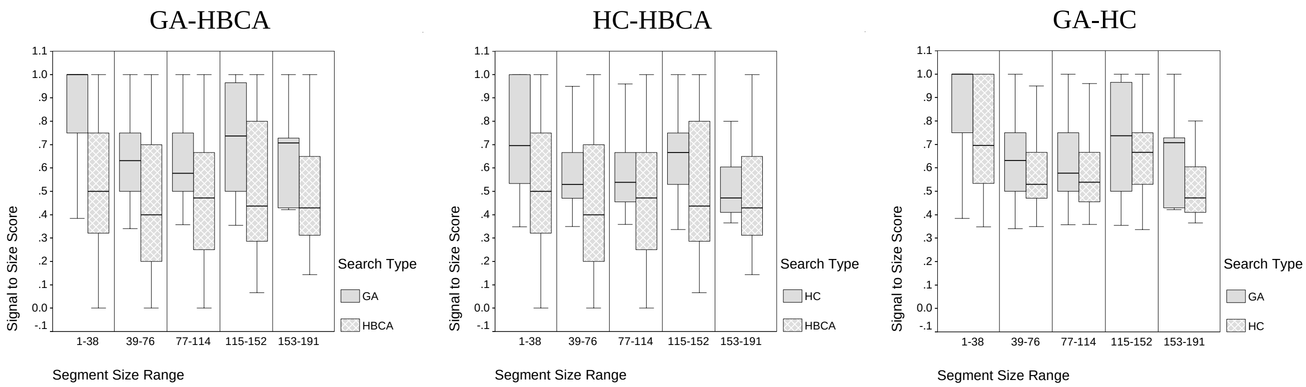

Figure 8. Boxplots of HBCA, GA and HC Signal to Size ratio compared across increasing Hypothesis List size range.

gram slicing [24] in conjunction with HBCA derived concepts to create executable concept slices(ECS) [7, 10]. ECS involves the slicing of concept bindings from the HBCA according to the system dependence graph approach of Horwitz et al [12]. The resulting ECS are proposed to posses the higher level abstraction of concepts alongside the useful executability of a program slice [7, 10].

Another recent related technique involves the use of Latent Semantic Analysis(LSA) for concept location [18]. "LSA is a fully automatic mathematical/statistical technique for extracting and inferring relations of expected contextual usage of words in passages of discourse [16]". The technique involves the analysis of user queries alongside the parsing and analysis of code as text to identify concepts.

6. Conclusions and Future Work

An approach to permit overlapping Concept Boundaries for Concept Assignment was presented in this paper. The problem was defined, analysed and formulated as a search problem in terms of search space, Fitness Function, GA and HC algorithms and solution structure in Section 3. An empirical study was carried out in two parts to determine the best algorithm in Section 4. First study compared the proposed GA, HC and Randomly generated solutions based on the proposed Fitness Function. The second study compared the GA and HC results with HBCA, based on Signal to Size ratio. The results of these studies were discussed in Subsections 4.1 and 4.2 respectively.

The GA results produced significantly stronger Fitness values according to the proposed Fitness Function. In addition, the GA results were significantly better than HC and HBCA results according to the Signal to Size ratios, as discussed in Subsection 4.2. This identified the GA as the best of the proposed algorithms for Concept Assignment which allow overlapping Concept Boundaries. On the other hand the HC results were somewhat disappointing as they were found to be significantly worst than GA and randomly generated solutions based on the proposed Fitness Function. However HC produced stronger results when compared to the HBCA on the Signal to Size measure. One possible explanation for this behaviour may be the increase in complexity of the search due to the inclusion of Hypothesis List Coverage as part of the fitness criteria, where a local search algorithm such as Hill Climbing is simply not adequate. Another explanation could be the inadequacy of the current neighbourhood definition and the need for examining alternative neighbourhood definitions. Another observation made was on the comparatively small range of HC Fitness values in Subsection 4.1 compared to other search algorithms. The smaller Fitness distribution implies that a set of similar Fitness values were achieved by the HC from random starting points, which in turn may indicate a large number of similar locally optimum solutions within the search space. Since the shape of the landscape is directly effected by the neighbourhood definition, this observation also strengthens the need for more suitable neighbourhood definitions for the HC algorithm.

Further research is required to analyse the resulting Concept Bindings as reflected in code. Useful future investigations may take the form of measuring the level of agreement for the location of Concept Bindings, analysis of the size and distributions of created Segments in the Hypothesis List and the size and distribution of the resulting Concept Bindings in program code. These investigation may also help to demonstrate the potential offered in program comprehension by the proposed techniques and whether the inclusion, extent and frequency of overlap could help or hinder program comprehension.

References

[1] T. J. Biggerstaff, B. G. Mittbander, and D. Webster. The concept assignment problem in program understanding. In *Proceedings of the 15th international conference on Software Engineering (ICSE 1993)*. IEEE Computer, 1993.

[2] P. Devanbu, R. Brachman, P. Selfridge, and B. Ballard. LaSSIE: A Knowledge-Based Software Information System. *CACM*, 34(5):34–49, May 1991.

[3] N. Gold. Hypothesis-Based Concept Assignment to Support Software Maintenance. In *ICSM*, pages 545–548, 2001.

[4] N. Gold and K. H. Bennett. A Flexible Method for Segmentation in Concept Assignmen. In *IWPC*, pages 135–144, 2001.

[5] N. Gold and K. H. Bennett. Hypothesis-based concept assignment in software maintenance. *IEE Proceedings - Software*, 149(4):103–110, 2002.

[6] N. E. Gold. *Hypothesis-Based Concept Assignment to Support Software Maintenance*. PhD thesis, Research Institute in Software Evolution, Department of Computer Science, University of Durham, 2000.

[7] N. E. Gold, M. Harman, D. Binkley, and R. M. Hierons. Unifying program slicing and concept assignment for higher-level executable source code extraction. *j-SPE*, 35(10):977–1006, Aug. 2005.

[8] D. E. Goldberg. *Genetic Algorithms in Search, Optimization and Machine Learning*. Addison-Wesley Publishing Company, 1989.

[9] D. E. Goldberg, B. Korb, and K. Deb. Messy genetic algorithms: Motivation, analysis, and first results. *Complex Systems*, 3:493–530, 1989.

[10] M. Harman, N. Gold, R. M. Hierons, and D. Binkley. Code Extraction Algorithms which Unify Slicing and Concept Assignment. In *WCRE*, pages 11–21, 2002.

[11] J. H. Holland. *Adaptation in Natural and Artificial Systems*. University of Michigan Press (Second edition MIT Press 1992), second edition, 1975.

[12] S. Horwitz, T. W. Reps, and D. Binkley. Interprocedural Slicing Using Dependence Graphs. In *PLDI*, pages 35–46, 1988.

[13] G. Huang and A. Lim. Designing a Hybrid Genetic Algorithm for the Linear Ordering Problem. In *Genetic and Evolutionary Computation - GECCO 2003*, pages 1053–1064. Springer, July 2003.

[14] V. Karakostas. Intelligent Search and Acquisition of Business Knowledge from Programs. *Journal of Software Maintenance: Research and Practice*, 4:1–17, 1992.

[15] D. L. Kreher and D. R. Stinson. *Combinatorial Algorithms. Generation, Enumeration and Search*. CRC Press, 1999.

[16] T. K. Landauer, P. W. Foltz, and D. Laham. An Introduction to Latent Semantic Analysis. *Discourse Processes*, 25:259–284, 1998.

[17] S. Letovsky and E. Soloway. Delocalized Plans and Program Comprehension. *IEEE Software*, 3(3):41–49, May 1986.

[18] A. Marcus, A. Sergeyev, V. Rajlich, and J. Maletic. An Information Retrieval Approach to Concept Location in Source Code. In *Proceedings of the 11th Working Conference on Reverse Engineering (WCRE 2004)*, pages 214–223, Nov. 2004.

[19] M. Mitchell. *An Introduction to Genetic Algorithms*. MIT Press, 7 edition, 2001.

[20] C. Rich and R. C. Waters. The programmer's apprentice. *Computer*, 21(11):10–25, Nov. 1988.

[21] E. Rodriguez-Tello and J. Torres-Jimenez. ERA: An Algorithm for Reducing the Epistasis of SAT Problems. In *Genetic and Evolutionary Computation - GECCO 2003*, pages 1283–1294. Springer, July 2003.

[22] P. Ross, J. G. Marine-Blazquez, S. Schulenburg, and E. Hart. Learning a Procedure That Can Solve Hard Bin-Packing Problems: A New GA-Based Approach to Hyper-heuristics. In *Genetic and Evolutionary Computation - GECCO 2003*, pages 1295–1306. Springer, July 2003.

[23] S. Tilley. A Reverse-Engineering Environment FrameWork. Technical Report CMU/SEI-98-TR-005, ESC-TR-98-005, Software Engineering Institute, Carnegie Mellon University, Pittsburgh, Pennsylvania 15213, April 1998.

[24] M. Weiser. Program slicing. *IEEE Transactions on Software Engineering*, 10:352–357, July 1984.

[25] L. Wills. *Automated Program Recognition by Graph Parsing*. PhD thesis, MIT Artificial Intelligence Lab, July 1992.