

## Research Article

# A Stochastic Hyperheuristic for Unsupervised Matching of Partial Information

**Kieran Greer**

*Distributed Computing Systems, Belfast, UK*

Correspondence should be addressed to Kieran Greer, [kgreer@distributedcomputingsystems.co.uk](mailto:kgreer@distributedcomputingsystems.co.uk)

Received 28 May 2012; Accepted 21 September 2012

Academic Editor: Thomas Mandl

Copyright © 2012 Kieran Greer. This is an open access article distributed under the Creative Commons Attribution License, which permits unrestricted use, distribution, and reproduction in any medium, provided the original work is properly cited.

This paper (Revised version of a white paper “Unsupervised Problem-Solving by Optimising through Comparisons,” originally published on DCS and Scribd, October 2011.) describes the implementation and functionality of a centralised problem solving system that is included as part of the distributed “licas” system. This is an open source framework for building service-based networks, similar to what you would do on a Cloud or SOA platform. While the framework can include autonomous and distributed behaviour, the problem-solving part can perform more complex centralised optimisation operations and then feed the results back into the network. The problem-solving system is based on a novel type of evaluation mechanism that prefers comparisons between solution results, over maximisation. This paper describes the advantages of that and gives some examples of where it might perform better, including possibilities related to a more cognitive system.

## 1. Introduction

This paper describes the implementation and functionality of a centralised problem-solving system that is included as part of the distributed “licas” service-based framework [1]. The licas (lightweight (internet-based) communication for autonomic services) system is an open source framework for building service-based networks, similar to what you would do on a Cloud or SOA platform. The framework comes with a server for running the services on, mechanisms for adding services to the server, mechanisms for linking services with each other, and mechanisms for allowing the services to communicate with each other. The default communication protocol inside of licas itself is an XML-RPC mechanism, but dynamic invocation of external Web Services is also possible. The main server package is now completely J2ME compatible, meaning that porting to a mobile device should be possible. The architecture and adaptive capabilities through dynamic linking add something new that is not available in other similar systems.

While the framework is built around distributed and autonomous objectives, the system is also useful as a test platform for more general AI problems. As such, a centralised component has been added, allowing for heuristic searches

to evaluate the situation and feedback the results. The centralised problem solver uses a hyperheuristic with a matching process at its core. The algorithm and novel nature of the process can be briefly described as follows. The solutions and the problem datasets are randomly placed into a grid and then a game is played to try and optimise the total cost over the whole grid. This is done by matching values across rows or columns. For the current problem only solution evaluations are matched, where to match any two solutions the algorithm must remove any rows that are in-between the two to be matched. The algorithm therefore also removes solutions as well as trying to keep other ones. While matching does not maximise, the algorithm tries to produce the largest overall score and therefore prefers to match higher valued solutions over lower valued ones. The philosophy behind the algorithm is described more completely in [2]. It would work particularly well for problems that might require some sort of symbolic evaluation instead of a numerical one. In that case, an exact evaluation of what value is “better” might not be possible and so some sort of matching evaluation would be required instead.

With regards to licas running information-based services, the purpose of the problem solver is to try to combine distributed sources of information through heuristic search

to generate more meaning over the information sources as a whole. The information to be combined can be partial in nature, or change over time. It is therefore difficult to evaluate accurately what pieces of information would belong together. Additional data could make an evaluation better or worse and so some sort of comparison or matching process might be preferred. The hyperheuristic is able to control this process and feed the results back into the network, so that the distributed information sources that are most likely to contain related information can be combined in an efficient and accurate manner. The problem solver itself can actually solve problems using either a hill-climbing approach or the new matching process, but this paper is concerned with the matching process only. This paper describes the new heuristic and considers the system integration details in particular. It also considers the scenarios where the new heuristic would work particularly well and describes how this differs from the more traditional clustering methods. The conclusions will also try to tie this in with a more cognitive model, which is something that the author is currently working on.

The rest of this paper is organised as follows. Section 2 summarises the main features of hyperheuristics and why they are useful. Section 3 discusses the problem of variable selection and some other heuristic algorithms that are commonly used for categorisation. Section 4 describes the new hyperheuristic framework in more detail, including implementation details. Section 5 describes some test results that show the heuristic working in an unsupervised manner. Section 6 gives some conclusions on the work and Section 7 describes future possibilities in the area of a more cognitive model.

## 2. Hyperheuristics

Most of this section has been taken from the literature review on hyperheuristics [3]. Many real-world problems that require some level of intelligence are difficult to solve. If all of the potential solutions can be realised, then the best one will be available and can be selected. Often, there are too many potential solutions and so heuristic search is required to estimate what the best solution might be. This is where intelligence is required, to help the heuristic to select what potential solutions should be explored further. The problem solving process is restricted to the information that is available in any potential solution. The search process can then reveal more information that was not originally known, but if the search space is very large, any solution will still only be an estimate or approximation of the true answer. The choice of heuristic that should be used to make this approximation then becomes very important, as different heuristics can evaluate certain concepts better than others. The main goal of hyperheuristics is to develop algorithms that are more generally applicable. Paper [4] is a recent survey of hyperheuristics. As noted in [5], a heuristic can be considered as a “rule of thumb” or “educated guess” that reduces the search required to find a solution. Allowing different types of evaluator to give a more complete picture of what the correct evaluation is. While a single heuristic can get stuck in locally optimal solutions, if several heuristics are compared, then

a more universal picture can be obtained. This can lead to better solutions somewhere else in the search space. The main drawback is that hyperheuristics need to be configured, or fine-tuned with the correct parameter settings, to work well. This is often a manual trial and error process.

The introductory sections in [4] make some interesting points. They note that hyperheuristics were initially developed as “heuristics to choose heuristics”. They are not intended to operate on the problem data itself, as a meta-heuristic would do. Instead, they operate on the heuristics that do evaluate the data directly, to select what solutions should be considered at each evaluation stage. With the incorporation of genetic programming, there is also the option of using “heuristics to generate heuristics”. The hyperheuristic can select which heuristics are mutated, or changed, for the next evaluation stage. The problem solving system of *licas* currently only uses heuristics to generate new heuristics, inside a genetic programming framework. Using heuristics to generate new heuristics not only involves selecting heuristics for the next evaluation stage, but also the ability to alter them resulting in a new heuristic not previously available. These sorts of problems can be solved by generating random solutions as part of a search process. Each solution is then changed in some way to improve it, until an optimal solution is obtained. The next stage of each search process is then directed by optimising the new solution set. Genetic programming itself is not inherently hyperheuristic, as it can also be used to represent the problem solutions. The hyperheuristic framework is more likely to use genetic programming principles to mutate existing solutions to generate better ones.

*2.1. Hyperheuristics Related to the New Problem Solver Heuristic.* This section looks at specific examples of heuristics that are directly related to the new hyperheuristic that is the core evaluator in the problem solver. The paper [6] describes a hyperheuristic framework that is self-organising by using reinforcement learning to order potential solutions. In this case, they apply each heuristic to a candidate solution to determine how it changes. If the solution changes positively, then the change is accepted. This is a perturbative approach, but includes both heuristic selection and heuristic creation or mutation. The low-level heuristics are evaluated through reinforcement learning into positive or negative ones. The positive ones are more likely to produce a positive evaluation and so are placed into a category of hill-climbing heuristics, to try to move directly to an optimal solution. The negative ones are placed into a category of mutational heuristics, which can then be changed to produce different ones.

Paper [7] would classify the new hyperheuristic as an evolutionary mechanism, because it contains a stochastic element and also allows for solution mutations. They note that while this can lead to mistrust in its use, it is also a more natural or bioinspired way of solving a problem. A key factor with this hyperheuristic, or for comparisons with other ones, is where in the process the randomness is applied. In most cases, the low-level heuristics can be changed in a random way, to generate new solutions that are then evaluated by the hyperheuristic for improvements to the current solution.

In this case, the randomness applies to the evaluation process of the hyperheuristic itself. Their own XCS algorithm uses the problem state to determine what heuristic to apply at some stage of the problem solving process. It also however, chooses randomly which problem to solve at each step, and randomness is also used as part of the problem-solving process itself. They also write.

*“The key idea in hyper-heuristics is to use members of a set of known and reasonably understood heuristics to transform the state of a problem. The key observation is a simple one: the strength of a heuristic often lies in its ability to make some good decisions on the route to fabricating an excellent solution. Why not, therefore, try to associate each heuristic with the problem conditions under which it flourishes and hence apply different heuristics to different parts or phases of the solution process? The alert reader will immediately notice an objection to this whole idea. Good decisions are not necessarily easily recognisable in isolation. It is a sequence of decisions that builds a solution, and so there can be considerable epistasis involved—that is, a non-linear interdependence between the parts. However, many general search procedures such as evolutionary algorithms and, in particular, classifier systems, can cope with a considerable degree of epistasis, so the objection is not necessarily fatal.”*

This appears to state that it is not always obvious or clear when a particular solution should be selected. As with nature, some level of randomness can be used to make an incorrect or imperfect selection process more robust. Paper [8] describes a hyperheuristic that also uses a simulated annealing approach for selecting which solutions to search further. The purpose of simulated annealing is also to add a stochastic element, to make the heuristic more generally applicable. The stochastic element can prevent a search from getting trapped in a local minimum, a place that a particular heuristic would naturally evaluate to, based on its limited knowledge. The stochastic element can help to make a decision that lies outside of the evaluation of the heuristic. The unpredictability however means that it can be hit-and-miss. Their algorithm adopts a simulated annealing acceptance criterion to alleviate the shortcomings of hill-climbing or exhaustive search. Their algorithm also uses stochastic heuristic selection mechanisms instead of deterministic ones, which has been shown to be superior for some evolutionary optimisation problems [9]. They evaluate a heuristic to get a score and then use simulated annealing to generate a probability threshold that the score must then match. The better solution score is more likely to meet the selection criteria. Paper [10] is also very interesting and tries to develop a hierarchical clustering algorithm that might be more applicable to the aims of the current project. It also describes other types of categorisation and matching functions not listed here.

### 3. Variable Selection

Before any entity can be analysed, it has to be determined what the most important features of that entity are. These features are then used to classify or evaluate the entity. This can be a difficult task because an entity could be composed of thousands of different features and so it is important to recognise the most important ones that make it different, or the same, as other entities. Paper [11] describes mechanisms for selecting variables or features from large repositories of unstructured data. These can act as filters, to select what variables or features from a potentially large dataset should be used to actually classify the dataset. It also notes that the most relevant variables are not necessarily the most useful when building a predictor or evaluator and so it is not simply a statistical matter of selecting the most popular variables. It is also possible to select subsets of variables that together have good predictive power. Papers [12, 13] discuss the difference between relevant and useful variables. In [12] they describe that at a conceptual level, one can divide the task of concept learning into two subtasks: deciding which features to use to describe the concept and deciding how to combine those features. The selection of relevant features, and the elimination of irrelevant ones, is one of the central problems in machine learning. Algorithms can range from something like nearest neighbour, which can calculate attribute distances based on all available information, to weighted feature selection, or even techniques for learning logical descriptions. The definition of relevance can mean [12] the following.

- (1) Relevant to the target concept.
- (2) Strongly or weakly relevant to a sample or distribution.
- (3) Relevant as a complexity measure.
- (4) Incremental usefulness.

Relevant to a target concept means that a change in the variable's value can change its classification allocated by the target concept. Relevant to a sample or distribution is the same, except for the fact that the variable is then required to be part of the sample, as well as relevant to the target concept. These notions are more important for an algorithm that is deciding which features to keep or ignore. Relevance to just the target concept can sometimes be used to try and prove the algorithm itself rather than its evaluating results. The new heuristic has potential for feature selection and would probably belong to category 2. In particular, for selecting the most appropriate values for certain variables or features from distributed or partial information. These evaluations can be better or worse than the true value and might vary around some distribution or mean of the true value. Each group of features or concepts can also be different in each solution part, but related or derived from a larger set. Sections 4 and 5 describe how the hyperheuristic can be used to try to select the best set of values for this type of scenario, in an unsupervised manner.

*3.1. Feature Selection Equations.* Existing feature selection usually involves categorising or clustering into distinct

groups. This is also often a supervised process, with known clusters being used to train the classifier, so that it can then recognise these clusters in other datasets as well. There are a number of existing equations that can be used to categorise data. Most of these actually belong to clustering algorithms that would try to measure how similar two individual data objects are, although some individual objects can be represented by cluster means. This is not actually what the matching process described in this paper is trying to do and so it already shows a possible difference in the use of the new hyperheuristic algorithm. Some of these equations are as follows.

**3.1.1. Euclidean Distance.** This is a linear measurement that is one of the simpler classification metrics. It sums the difference between all attributes of two different input objects to determine how similar they are to each other. The equation can look like [14]:

$$d_{12} = \sum_{j=1}^k (\pi_{1j} - \pi_{2j})^2, \quad (1)$$

where  $d$  is the distance and  $\pi_1$  or  $\pi_2$  are the input objects.

**3.1.2. Kullback-Leibler Information Divergence.** The Kullback-Leibler information divergence is a measure of the difference between two probability distributions. As described in Wikipedia: it can be used as a distance metric and measures the expected number of extra bits required to code samples from  $P$  when using a code based on  $Q$ , rather than using a code based on  $P$ . The equation can look like [14]:

$$D(p||q) = \sum_{j=1}^n p_j \log\left(\frac{p_j}{q_j}\right) \quad \text{for } p = (p_1, p_2, \dots, p_n), \\ q = (q_1, q_2, \dots, q_n). \quad (2)$$

This measures distances between probability distributions, instead of single dataset values and is therefore probably more useful for measuring the distance between created cluster groups, than the individual data objects.

**3.1.3. Jaccard Coefficient.** The Jaccard coefficient measures the similarity between datasets. It is a set theoretic measure and can be defined as the intersection of the datasets divided by the union of the datasets. For example,

the Jaccard coefficient between  $T_1$  and  $T_2$  can be defined as  $|(T_1 \cap T_2)/(T_1 \cup T_2)|$ .

Dividing by the intersection scales the result between 0 and 1. If the two sets are the same, for example, the equation computes to the value 1. If there are no elements the same, then it computes to 0. The Jaccard distance measure is then the opposite of this and measures the dissimilarity between two datasets. One drawback of the Jaccard coefficient is that it does not really consider negative input as well.

**3.1.4. Rocchio Classifier.** The Rocchio classifier [15] is a similarity-based linear classifier that considers both positive and negative input. Equation [16] can be described as given a training dataset  $T_r$ , the Rocchio classifier directly computes a classifier  $\bar{c}_i = \langle w_{1i}, w_{2i}, \dots, w_{ri} \rangle$  for category  $c_i$  by means of the formula:

$$w_{ki} = \beta \cdot \sum_{\{d_j \in \text{POS}_i\}} \frac{w_{kj}}{|\text{POS}_i|} - \gamma \cdot \sum_{\{d_j \in \text{NEG}_i\}} \frac{w_{kj}}{|\text{NEG}_i|}, \quad (3)$$

where  $w_{kj}$  is the weight of dataset  $t_k$  in document  $d_j$ , and POS or NEG means that document  $d_j$  contained in the training dataset, does or does not belong to the classifier category  $c_i$ . A classifier built using the Rocchio method rewards the closeness of a test document to the centroid of the positive training examples and its distance from the centroid of the negative training examples.

**3.1.5. Information Theoretic Ranking Using Probability Densities.** Paper [11] gives an example of a ranking equation that can be used with information theoretic criteria. Information theory has to do with data compression and also loss of information through noise. The following is an example of the sort of equation that would be used to evaluate that. This ranking equation can be used to determine the probability of one variable being associated with another one, or some target concept. This relies on probability densities that can be unknown or hard to estimate. With discrete or nominal variables however, it can be written as

$$I(i) = \sum_{x_i} \sum_y P(X = x_i, Y = y) \log \frac{P(X = x_i, Y = y)}{P(X = x_i)P(Y = y)}, \quad (4)$$

where  $P(x_i)$  is the probability density for variable  $x$  at time  $i$ ,  $P(y)$  is the probability density for target  $y$ , and  $P(x, y)$  is the probability of them occurring together. The value is therefore a measure of the dependency between the target and the variable in question.

**3.1.6. Information Gain.** Paper [17] also describes information gain. With this method, both class membership and the presence/absence of a particular term are seen as random variables, and one computes how much information about the class membership is gained by knowing the presence/absence. Indeed, if the class membership is interpreted as a random variable  $C$  with two values, positive and negative, and a word is likewise seen as a random variable  $T$  with two values, present and absent, then using the information-theoretic definition of mutual information we may define Information Gain as

$$\text{IG}(t) = H(C) - H(C | T) \\ = \sum_{c, t} P(C = c, T = t) \ln \left[ \frac{P(C = c, T = t)}{P(C = c)P(T = t)} \right]. \quad (5)$$

Here,  $t$  ranges over present, absent and  $c$  ranges over  $\{c+, c-\}$ . As pointed out above, this is the amount of information about  $C$  (the class label) gained by knowing  $T$  (the presence or absence of a given word).

3.1.7. *Feature Selection Based on Linear Classifiers.* As described in [17], both SVM and Perceptron, when used as linear classifiers, output predictions of the form:

$$\text{prediction}(x) = \text{sgn}(w^T x + b) = \text{sgn}(S_j w_j x_j + b). \quad (6)$$

Thus, a feature  $j$  with the weight  $w_j$  close to 0 has a smaller effect on the prediction than features with large absolute values of  $w_j$ . The weight vector  $w$  can also be seen as the normal to the hyperplane determined by the classifier, to separate positive from negative instances. Thus we often refer to the procedure as “normal based feature selection”. One speculates that since features with small  $|w_j|$  are not important for categorization they may also not be important for learning and, therefore, are good candidates for removal. A theoretical justification for retaining the highest weighted features is to consider the feature important if it significantly influences the width of the margin of the resulting hyperplane. This is described further in the paper.

#### 4. New Stochastic Hyperheuristic Framework

The new hyperheuristic framework was first introduced in [2]. Essentially, it uses a matching evaluation over a maximising one. However, it also tries to maximise the matching score and so will favour higher scoring matches over lower scoring ones. It also uses a randomising procedure to place all potential solutions in a grid, where any solution can be placed in any position. The matching process then matches solutions by removing any solutions that are between them in the grid and grouping the matched ones together for evolving into a new solution. This process therefore also removes solutions as well as evolving the potentially better ones and in that sense is self-regulating. There are, however, other parameters that need to be set during configuration.

This sort of process might be preferable for the feature selection problem that has been described, or the heuristic might, in general, be more suitable for a different class or type of problem. The clustering heuristics that have been described in Section 3 are intended to categorise similar datasets through a matching process, where the datasets with the most similar characteristics are grouped together. This therefore requires several category types and then several datasets belonging to each category type. The aim of the new heuristic is to try to “optimise some global evaluation” through a similar matching process. This evaluation applies to the data or problem set more as a whole, rather than evaluating each solution as a separate entity. It is more similar to a neural network trying to realise a single evaluation function that maps its input to its output, than several evaluation functions, each mapping a different value set. It is not a case of categorising the different datasets individually into similar groups, but rather, trying to evolve all solutions in the most robust way, in order to arrive at an optimal collective solution value. The whole search space belongs to the same single problem. Some of the best potential solutions can be removed, if it means that other ones match better as a result. Each matching phase is probably also associated with an evolution of the related solutions, to produce offspring that

would then more closely match or solve the problem. The correct evolutions are not known beforehand and so it is consistency through matching that is used to decide which solutions to evolve. The idea is that a more robust solution pool will be created, even at the cost of some very good ones and this will lead the search to a better global optimum, when keeping locally good solutions might lead to getting stuck in a local optimum.

4.1. *Implementation Details.* The problem solving framework has now been implemented as part of the *licas* system [1]. *Licas* provides a framework for building distributed service-based networks of information sources, for example. The individual services can self-organise through a novel linking mechanism and can also display autonomous behaviours. The self-organising mechanism is relatively lightweight and essentially stores links from one service or node to another represented by a weight strength. The links are made more accurate through a path description made up of metadata or concepts that relate to the association between the linked services or nodes. So there is not a great deal of computation that takes place to create these links. It is also a highly distributed solution, where one link or association does not have to relate to any other one. It is built up purely through the feedback of the system use and does not use any centralised or knowledgeable algorithm.

The problem solver is then more of a centralised solution. It can be sent the information from the network sources, use heuristic search and evaluations to perform a more complex problem-solving operation and then feed these results back into the network, to allow the sources to update themselves through the more complex search procedure. The problem solver can also be used by itself without the network, to solve any sort of problem using genetic algorithms, where the framework is very extendible with the user’s own classes. Figure 1 is a schematic of the general problem solving framework. The problem solving is performed locally and not distributed throughout all of the network services or nodes.

Any test problem can be configured using a test script. The *licas* system also has a GUI that can be used to run the tests. An information mediator can be used to send and retrieve the information from the distributed sources. Each source can be used to create a service that is initialised with the information and then run on the network. The paths to the data sources can be specified in the script-currently file paths. The mediator then periodically asks the services for their recent evaluations or information and invokes the problem solver to cluster or solve the information set as best it can. The result of this can then be sent back to the services. It is currently turned into dynamic links, to update the network structure.

#### 5. Testing

A set of tests have been created to test the problem solving framework for usefulness as a feature selector. The tests are designed mainly to determine if the problem-solving process

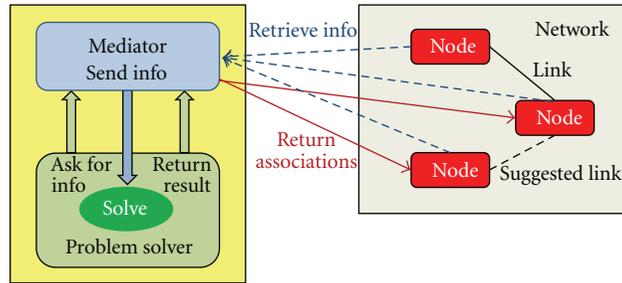


FIGURE 1: General problem-solving framework.

is constructive, that is, it is doing something in an intelligent manner and not simply trying to match things in a random way. The test was as follows: datasets of concepts or features were created. The features were just letters or letter combinations of the alphabet, where each dataset was assigned a random selection of these features. Each dataset therefore had a randomly selected subset of all of the features. Each feature was assigned a random score, but distributed around some mean value. There was a preferred or more common mean value, with a distribution of values either side of it. Each feature therefore had a random score, but more commonly, this score would be closer to the mean value of the distribution. The matching process therefore should match the more common scores and therefore choose those related solutions for evolving further. It would also naturally select certain groups of features for evolving as part of the process.

The evaluation and evolution process is simply an intersection of the dataset features. During the comparison, if both datasets have certain features in common, their distribution from the mean for those features only is calculated and this is taken to be the similarity score. It is also only those features that are then included in an evolved solution. The evolution algorithm evolves using an intersection and so the sums for single datasets will generally be larger than for evolved sets. A hill-climbing approach did not work as well for this type of problem, because maximising the count would prefer the original datasets with all of their features and largest values, over an evolved one with only subsets of those features. The hyperheuristic using a similarity approach, preferred combining datasets with the smallest difference in score, when the evaluation score could be any size. Other evaluation equations were also tried, such as Euclidean, Jaccard, or similarity, see Section 3.1, but the intersection of features proved to be a more suitable approach.

*5.1. Example Test Data.* One example test was as follows: 30 datasets were created randomly. Each dataset contained 20 random features selected from a possible 40 features. Each feature was assigned a random score taken from the distribution described in Table 1. This description shows that there was a mean value of 10 that would be selected 40% of the time. This is indicated by the 0.4 percentage value and the 1.0 distribution value in the centre of the script description.

There was then a distribution, either of values 11–13 or 7–9 that would be selected 17.5% of the time. Either side of the centre value are two sets with 0.175 percentages and then a distribution change of  $\pm 3.0$ . The next distribution was 14–16 or 4–6, selected 7.5% of the time. Finally, the distribution of either the values 17–19 or 1–3 would be selected 5% of the time.

A random process might be expected to produce an average, or 50% deviation, from the true mean value of 10, when selecting variable or feature values. If considering just the difference from the mean value, if 10 separate values are selected then this could be calculated on average as the mean value 10 is selected 4 times, 8 is selected 3.5 times, 5 is selected 1.5 times, and 2 is selected 1 time. Note that 8 selected 3.5 times relates to doubling up on either 8 or 12 (both sides), selected 1.75 times each, and so forth. The difference is being measured here and not the exact total. This gives a total average of

$$40 + 28 + 7.5 + 2 = \frac{77.5}{10} = 7.75. \quad (7)$$

This is a difference of 2.25 from the true mean value of 10. The success of the test was then calculated as follows. Some of the final clusters contained the original solution datasets, while others contained the evolved solution datasets. If the process is constructive or intelligent, then the evolved solutions should be closer to the mean value than the original ones. Therefore the stats calculated the difference to the mean for the original solutions that were part of the final solution set and also the evolved solutions that were part of the final solution set. The original solutions that were included actually had an average difference of possibly around 2.27, which is close to the random average value. The evolved solutions had an average difference of possibly 1.57, which is a 30% improvement on the random or single datasets value. A 30% improvement in something might be good or bad, depending on the particular application or problem. The test data, while random, is not from a real application. So while this might not be able to prove a better approach, it shows that the process is constructive and is doing something intelligent. It is unsupervised and therefore all of the evaluations are taken from the information that is presented at the time. Using the intersection method as the evaluator and evolver is again just for initial results, where further tests could find some improvement there as well.

TABLE 1: Distribution values used in the tests.

Percentage distribution	5%	7.5%	17.5%	40%	17.5%	7.5%	5%
Value distribution	1–3	4–6	7–9	10	11–13	14–16	17–19
Program script	0.05 : 3	0.075 : 3	0.175 : 3	0.4 : 1	0.175 : 3	0.075 : 3	0.05 : 3

- (1) Read the variable set and the distribution values.
- (2) Generate a number of datasets from the set of variables.
  - (a) Randomly select variables from the set to add to the dataset.
  - (b) Randomly select a count from the distribution, for each variable.
- (3) Save the test dataset with a unique name.

ALGORITHM 1: Algorithm for generating the test data.

- (1) Generate a number of datasets from the spec and random generator.
- (2) Read each dataset and generate a test service (and solution) from it.
- (3) Evaluate the fitness, or value, of each solution.
- (4) Place the solution evaluations in the grid, in a random order, and try to match.
  - (a) A match is determined by a pre-defined allowed amount of difference.
  - (b) Only a specified number of final matches are kept.
  - (c) Also, keep the matches with the larger values.
  - (d) Matched solutions are stored as pairs to be evolved.
- (5) Evolve the matched solutions to produce a new test set of original plus matched solutions.
- (6) The evolved solution can be an intersection or mutation of the original two solutions.
- (7) Measure the total average variable value of existing original solutions in the final set and the total average variable value of evolved solutions in the final set.
- (8) The evolved solutions should have an average value closer to the mean.

ALGORITHM 2: Algorithm for evaluating the test data.

5.2. *Test Algorithm.* A more complete description of how the new hyper-heuristic framework works is given in [2]. This section includes algorithms specifically for the test that was carried out. The test data was created from Algorithm 1.

The datasets were evaluated and solutions generated based on the Algorithm 2.

## 6. Conclusions

This paper has described a new type of problem-solving framework, based more on corroborative or matching evidence, than on purely maximising some function. The aims of the current project are to develop a hyperheuristic framework for evaluating information sources, to try to combine sources that might be related. The problem itself however might not be a typical categorisation one, but one that tries to realise some global function over the whole set of problem data. It might be more similar to a neural network trying to realise a single evaluation function than several evaluation functions that each map a different thing. Tests have shown that it could be useful, for example, for selecting feature groups or best values out of partial information sets. Each dataset contains only parts of the whole solution or picture, when combining and evaluating these parts produce new solutions and parts that are closer to the whole picture.

While the data was random, it would still contain certain patterns, based on the created feature groups. This would be strengthened by the count and count additions, associated with each feature and new solution.

The work arose from looking at the problem of aggregating distributed information sources or concepts autonomously, to try to formulate more complex real-world entities. For example, can a sensorised system that is fed information from distributed sources, determine for itself, what concepts it receives are key and belong together as a more complete entity. If the underlying framework is unintelligent, relying more on statistical updates, then the stochastic element could help to provide more robust solutions. Also, because there is no inherent intelligence, there should be no bias towards any particular solution, where the statistical process should determine this for itself.

## 7. Future Work

The author also has an interest in developing a more brain-like or cognitive model. While the distributed part of this model has been written about recently [18], there is also possibly a centralised part to the model as well. This has been noted and could even be thought of as the consciousness part. If the neurons in the brain form in a purely mechanical

way, without any real intelligence, then a more centralised and intelligent part that can interpret the firing neurons might be required. It is also known that the human brain does not store information only once, but duplicates it in different areas. When the neurons fire, either a stronger signal from one set of neurons, or possibly from duplication of the firing pattern in different areas, could determine what the correct thought or decision is. If several areas fire at the same time with parts related to the same problem, this would produce a stronger overall signal and could help to suggest what the correct interpretation is. This is especially useful if the search process is not particularly structured or constructive, that is without a clear search path, for example, when the duplication can help to confirm what the correct interpretation is. This also means that a matching process is more attractive as the decision maker. If the decision is based on a signal strength from one place only, then the actual measurement of this has to be interpreted slightly more accurately and the neuron would need to do more. Although, more inputs to the neuron would also simply change this. If the signal is stronger through duplication however, then a more simple matching process can possibly derive the same conclusion. The author therefore finds the hyperheuristic that uses comparisons more attractive as the centralised component for his current cognitive model. A decision would be formed through more neuronal areas firing relevant signals when the input is received.

## References

- [1] Licas, <http://licas.sourceforge.net/>.
- [2] K. Greer, "A stochastic hyper-heuristic for optimising through comparisons," in *Proceedings of the 3rd International Symposium on Knowledge Acquisition and Modeling (KAM'10)*, pp. 325–328, IEEE, Wuhan, China, October 2010.
- [3] K. Greer, "Literature review for the multi-source intelligence project called "a stochastic hyper-heuristic for optimising through comparisons"" Distributed Computing Systems Research Report, 2011, <http://www.scribd.com/doc/58227009/Hyper-Heuristic-Literature-Review>.
- [4] E. K. Burke, M. Hyde, G. Kendall, G. Ochoa, E. Ozcan, and R. Qu, "A survey of hyper-heuristics," Computer Science Technical Report NOTTCS-TR-SUB-0906241418-2747, University of Nottingham, 2009, (hhSurvey09).
- [5] M. Bader-El-Den and R. Poli, "Generating SAT Local-Search Heuristics using a GP Hyper-Heuristic Framework," in *Proceedings of the 8th International Conference on Artificial Evolution (EA'07)*, pp. 37–49, 2007.
- [6] E. Ozcan, M. Misir, and E. K. Burke, "A self-organising hyper-heuristic framework," in *Proceedings of the 4th Multidisciplinary International Scheduling Conference: Theory & Applications (MISTA'09)*, pp. 784–787, Dublin, Ireland, August 2009.
- [7] J. G. Marín-Blázquez and S. Schulenburg, "Multi-step environment learning classifier systems applied to hyper-heuristics," in *Proceedings of the 8th Annual Genetic and Evolutionary Computation Conference*, pp. 1521–1528, Washington, DC, USA, July 2006.
- [8] R. Bai, J. Blazewicz, E. K. Burke, G. Kendall, and B. McCollum, "A simulated annealing hyper-heuristic methodology for flexible decision support," Tech. Rep. NOTTCS-TR-2007-8, School of CSiT, University of Nottingham, 2007.
- [9] T. P. Runarsson and X. Yao, "Stochastic ranking for constrained evolutionary optimization," *IEEE Transactions on Evolutionary Computation*, vol. 4, no. 3, pp. 284–294, 2000.
- [10] N. Sahoo, J. Callan, R. Krishnan, G. Duncan, and R. Padman, "Incremental hierarchical clustering of text documents," in *Proceedings of the 15th ACM Conference on Information and Knowledge Management (CIKM'06)*, pp. 357–366, New York, NY, USA, November 2006.
- [11] I. Guyon and A. Elisseeff, "An introduction to variable and feature selection," *Journal of Machine Learning Research*, vol. 3, pp. 1157–1182, 1993.
- [12] A. L. Blum and P. Langley, "Selection of relevant features and examples in machine learning," *Artificial Intelligence*, vol. 97, no. 1-2, pp. 245–271, 1997.
- [13] R. Kohavi and G. H. John, "Wrappers for feature subset selection," *Artificial Intelligence*, vol. 97, no. 1-2, pp. 273–324, 1997.
- [14] S. McClean, B. Scotney, K. Greer, and R. Pairceir, "Conceptual clustering of heterogeneous distributed databases," in *Proceedings of the 12th Joint European Conference on Machine Learning (ECML'01) and 5th European Conference on Principles and Practice of Knowledge Discovery in Databases (PKDD'01), Workshop on Ubiquitous Data Mining for Mobile and Distributed Environments*, pp. 46–55, September 2001.
- [15] J. J. Rocchio, "Relevance feedback in information retrieval," in *The SMART Retrieval System: Experiments in Automatic Document Processing*, G. Salton, Ed., pp. 313–323, Prentice Hall, Englewood Cliffs, NJ, USA, 1971.
- [16] G. Guo, H. Wang, D. Bell, Y. Bi, and K. Greer, "An kNN model-based approach and its application in text categorization," in *Computational Linguistics and Intelligent Text Processing, 5th International Conference, Ciling 2004, Seoul, Korea*, A. Gelbukh, Ed., pp. 559–570, Springer, New York, NY, USA, 2004.
- [17] D. Mladenić, J. Brank, M. Grobelnik, and N. Milic-Frayling, "Feature selection using linear classifier weights: Interaction with classification models," in *Proceedings of Sheffield SIGIR—27th Annual International ACM SIGIR Conference on Research and Development in Information Retrieval*, pp. 234–241, Sheffield, UK, July 2004.
- [18] K. Greer, "Symbolic neural networks for clustering higher-level concepts," *NAUN International Journal of Computers*, vol. 5, no. 3, pp. 378–386, 2011.



**Hindawi**

Submit your manuscripts at  
<http://www.hindawi.com>

