

A Proposition on Memes and Meta-Memes in Computing for Higher-Order Learning

Ryan Meuth¹, Meng-Hiot Lim², Yew-Soon Ong³, Donald C. Wunsch II¹

¹Applied Computational Intelligence Laboratory, Department of Electrical and Computer Engineering, Missouri University of Science and Technology, Rolla, MO, USA.
rmeuth@mst.edu

²School of Electrical & Electronic Engineering, Nanyang Technological University, Singapore 639798.

³School of Computer Engineering, Nanyang Technological University, Singapore 639798.

Keywords: machine learning, memetic computing, meta-learning, computational intelligence architectures

Abstract

In computational intelligence, the term ‘memetic algorithm’ has come to be associated with the algorithmic pairing of a global search method with a local search method. In a sociological context, a ‘meme’ has been loosely defined as a unit of cultural information, the social analog of genes for individuals. Both of these definitions are inadequate, as ‘memetic algorithm’ is too specific, and ultimately a misnomer, as much as a ‘meme’ is defined too generally to be of scientific use. In this paper, we extend the notion of memes from a computational viewpoint and explore the purpose, definitions, design guidelines and architecture for effective memetic computing. Utilizing two conceptual case studies, we illustrate the power of high-order meme-based learning. With applications ranging from cognitive science to machine learning, memetic computing has the potential to provide much-needed stimulation to the field of computational intelligence by providing a framework for higher order learning.

1. Introduction

Over the past several years many hundreds of papers have been published on the modification and application of only a handful of core computational intelligence techniques – namely dynamic programming, evolutionary algorithms, neural networks, fuzzy logic, and data clustering methods. Algorithmically, there have been refinements and crossovers in these categories, such as heuristic dynamic programming, particle swarm optimization, evolutionary-trained fuzzy neural networks, and hybrid genetic algorithms, resulting in significant but relatively modest quality and performance gains. Beyond these modifications the pace of new algorithm design has been stagnant for a period of time, while the complexity of machine learning and optimization problems has grown ever larger with the maturity of the internet, digital media, and the proliferation of data sources in all aspects of human life.

Meanwhile, advancement in hardware technology has brought about affordable and powerful computing platforms which are more easily accessible. However, it is clear that increase in computational capacity cannot even come close to addressing the challenges posed by the complexity of problems, many of which are typical of real-world scenarios (Gaudiot, Kang et al. 2005). More advanced and novel computational paradigms, particularly from the algorithms front have to be championed. The general perception on how algorithms have managed to keep pace with increasing problem complexity over the years is depicted in Figure 1. Initially, algorithms by and large were able to keep up with the demands of increasing problem complexity. To a certain extent, the algorithms which typically belong to the category of conventional or exact enumerative procedures were able to surpass the complexity of problems that were typical of what people were trying to solve. Subsequently, as the complexity of problems pushes the capability limits of algorithms, it became evident that the complexity of problems being addressed began to overwhelm the algorithms available. We view the region corresponding to the convergence and divergence of the curves as being synonymous to the era of computational intelligence techniques. It can be envisaged that in time, the spread between complexity of problems and algorithms will widen if computational intelligence remains at status quo. There are clear signs that these issues are in the early stages of being addressed. In particular, the phase of research should be putting

emphasis not just on learning *per se*, but rather on issues pertaining to higher order learning. This is a natural tendency in order to address the demands and challenges of problems that surface.

The era of computational intelligence to a certain extent managed to contain the gap between algorithms and problem. In time, it will become clear that the divergence between the two curves will continue, as shown in Figure 1. A more promising outlook as shown by the broken line curve can be achieved and modern day optimization techniques can rise to this challenge by incorporating not just mechanisms for adaptation during the process of solving an instance of a difficult problem, but rather mechanisms of adaptation or more appropriately learning spanning across instances of problems encountered during the course of optimization. While a certain degree of similarity may be drawn when compared to case-based reasoning (CBR), such perceived “experiential” trait similarity in the sense that both encompass mechanisms to draw on “experience” from previously encountered problem instances is superficial. Unlike CBR methods which rely on the need for explicit examples and ranking procedures, optimization problems are usually not amenable to such explicit case by case assessment to yield information that is potentially useful to a search algorithm (Kolodner 1993; Wills and Kolodner 1994). Rather, a more likely emphasis should be the building up of a body of knowledge, more specifically memes and meta-memes that collectively offer capability with a much broader problem-solving scope in order to deal with the class of problems being addressed.

In 1997 Wolpert and Macready formalized the “No Free Lunch Theorem” stating simply:

“Any two [optimization] algorithms are equivalent when their performance is averaged across all possible problems.”

Additionally, Wolpert and Macready made the observation that in order to reduce the average cost across a set of problems and optimizers, one must methodically utilize prior or acquired information about the matching of problems to procedures, given *a priori* knowledge gained from experience (Wolpert and Macready 1997). The realizations brought by the No Free Lunch Theorem changed the research focus of the field of computational intelligence from the design of individual algorithms to the design of architectures of algorithms and parameters optimization. It is in this spirit that the development of memetic algorithms has been motivated (Xin 1999; Kuncheva and Jain 2000; Francois and Lavergne 2001; Kazarlis, Papadakis et al. 2001; Lee, Lau et al. 2001; Milano and Roli 2004; Ong and Keane 2004; Krasnogor and Smith 2005; Ong, Lim et al. 2006).

Taken alone, current methods tend to be overwhelmed by large datasets and suffer from the curse of dimensionality. A new class of higher order learning algorithms are needed that can autonomously discern patterns in data that exist on multiple temporal and spatial scales, and across multiple modes of input. These new algorithms can be architectures utilizing existing methods as elements, but to design these architectures effectively, some design principles should be explored.

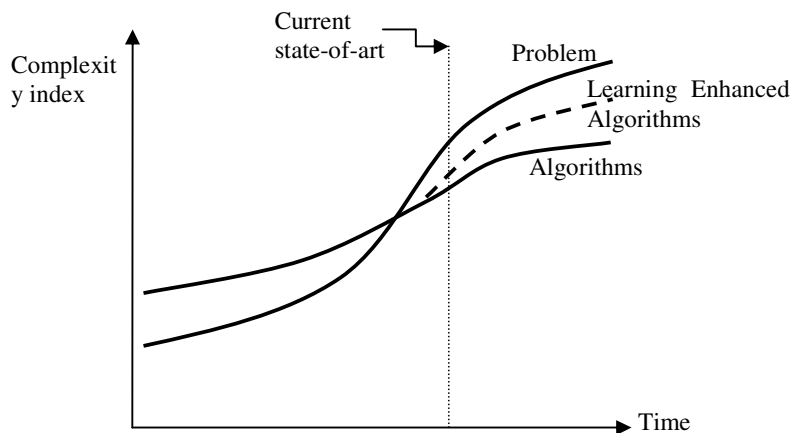


Figure 1: An abstract comparison on state of optimization from the perspectives of problems and algorithms complexity.

Ultimately, the curse of complexity cannot be wholly avoided. As the size or dimension of the problems increases, a greater amount of computation becomes necessary to find high quality solutions. However, such computation need not be done on the fly, meaning at the exact time that a problem is presented. If a memory

mechanism is provided that can store and retrieve previously used or generalized solutions, then computation can be shifted into the past, greatly reducing the amount of computation necessary to arrive at a high quality solution at the time of problem presentation.

One of the major drawbacks of evolutionary algorithms and computational intelligence methods in general is the solvers employed usually start from zero information, independent of how similar the problem instance is to other instances the method has been applied to in the past. In effect, the optimization methods typically do not incorporate any mechanisms to establish inter-instance memory. This property is useful for comparing different computational intelligence methods and in some cases, particularly when computation time is not an issue, the capacity to draw on memory of past instances solved is desirable as it allows the search to be more focused, thus leading to solutions that would not otherwise have been found efficiently. It is also worth noting that many real-world problem domains are composed of sub-problems that can be solved individually, and combined (often in a non-trivial way) to provide a solution for the larger problem (Lenat and Guha 1989; Shahaf and Amir 2007).

In some problem instances, such as large instances of the even parity problem, it is nearly impossible to stochastically arrive at a complete solution without utilizing generalized solutions for small instances of the problem (Koza 1989). It is simple to evolve a function that performs even parity on 2 bits using only the logical functions AND, OR and NOT as primitives, but extremely difficult to evolve a 10-bit even parity function without any *a priori* information as the space of all possible solutions is immensely larger, and even the best known solution is complex. By simply defining the general 2-bit XOR function (the even parity computation for 2 bits), the optimization method has a higher probability of combining instances of XOR to arrive at an n-bit even-parity function, greatly accelerating the optimization process.

In the game of chess, humans start at the top, and solve a successive sequence of smaller, tractable problems to arrive at a move. However, the learning process is bottom-up - a human player of chess first learns the legal moves of every piece, and then combines those general move capabilities into strategies, strategies into tactics and those tactics combine with the tactics of the opposing player to form a high-level view of the game as a whole. At each level optimization and generalization are performed to pass information up and down the play hierarchy. However, this natural progression is not reflected in the methods that we utilize to computationally approach problems of this scale. The typical approach is combinatorial optimization, where a sequence of low-level moves is statistically analyzed in order to arrive at a plan of play. As a whole, this is a computationally intractable problem, and it does not even come close to resembling the way humans play chess. Additionally, the skills learned in chess may translate across several domains as general problem solving skills. The ability to translate knowledge from one domain to another implies the necessity of meta-learning or learning about how or what to learn - in order to recognize similar problem features in disparate environments and scenarios.

The remaining of this paper is organized as follows. Section 2 gives a brief outline of the classes of brain inspired memetic computing. In Section 3 we discuss and compare between schema and memes, in particular their roles in learning. Section 4 gives an architectural framework for computing with memes and meta-memes, exposing some important issues in the design of systems with higher order learning capability. Two examples, the even parity in Section 5 and travelling salesman problem in Section 6 are studied to illustrate the concept of learning that spans across instances of problems. In Section 7, we conclude this paper.

2. Brain Inspired Memetic Computing

While Darwinian evolution has been a source of inspiration for a class of algorithms for problem-solving, memetics has served as a motivation for problem-solving techniques with memetic algorithms being the most prominent and direct manifestation of the inspiration. In recent years, there has been a marked increase in research interests and activities in the field of Memetic Algorithms. The first generation of MA refers to hybrid algorithms, a marriage between a population-based global search (often in the form of an evolutionary algorithm) coupled with a cultural evolutionary stage. The first generation of MA though it encompasses characteristics of cultural evolution (in the form of local refinement) in the search cycle, may not qualify as a true evolving system according to Universal Darwinism, since all the core principles of inheritance/memetic transmission, variation and selection are missing. This suggests why the term MA stirs up criticisms and controversies among researchers when first introduced in (Moscato 1989). The typical design issues (Nguyen, Ong et al. 2007) include i) how often should individual learning be applied, ii) on which solutions should individual learning be used, iii) how long should individual learning be run, iv) what maximum computational budget to allocate for

individual learning, and v) what individual learning method or meme should be used for a particular problem, sub-problem or individual.

Multi-meme (Krasnogor 2002), hyper-heuristic (Kendall, Soubeiga et al. 2002) and meta-Lamarckian MA (Ong and Keane 2004; Ong, Lim et al. 2006) are referred to as second generation MA exhibiting the principles of memetic transmission and selection in their design (Nguyen, Ong et al. 2008). In multi-meme MA, the memetic material is encoded as part of the genotype. Subsequently, the decoded meme of each respective individual is then used to perform a local refinement. The memetic material is then transmitted through a simple inheritance mechanism from parent to offspring. On the other hand, in hyper-heuristic and meta-Lamarckian MA, the pool of candidate memes considered will compete, based on their past merits in generating local improvements through a reward mechanism, deciding on which meme to be selected to proceed for future local refinements. A meme having higher rewards will have greater chances of being replicated or copied subsequently. For a review on second generation MA, i.e., MA considering multiple individual learning methods within an evolutionary system, the reader is referred to (Ong, Lim et al. 2006). Co-evolution and self-generation MAs introduced in (Krasnogor and Gustafson 2004) and (Smith 2007) are described in (Nguyen, Ong et al. 2008) as 3rd generation MA where all three principles satisfying the definitions of a basic evolving system has been considered. In contrast to 2nd generation MA which assumes the pool of memes to be used being known *a priori*, a rule-based representation of local search is co-adapted alongside candidate solutions within the evolutionary system, thus capturing regular repeated features or patterns in the problem space.

From the 3 classes of MA outlined, memes can be seen as mechanisms that capture the essence of knowledge in the form of procedures that affect the transition of solutions during a search. The level of participation or activation of memes is typically dictated by certain indicative performance metrics, the objective being to achieve a healthy balance between local and global search. Memes instead of being performance-driven should be extended to include capacity to evolve based on the snapshots of problem instances. In the process of solving a repertoire of problem instances, memes can culminate based on the recurrence of patterns or structures. From basic patterns or structures, more complex higher level structures can arise. In this regard, a brain inspired meta-learning memetic computational system, consisting of an optimizer, a memory, a selection mechanism, and a generalization mechanism that conceptualizes memes not just within the scope of a problem instance, but rather in a more generic contextual scope is appropriate. Such traits which are lacking in the 3rd generation MA can serve as the basis of 4th generation class of MAs. The reader is referred to Table 1 for a summary of generational description of Memetic Algorithms. The summary although by no means exhaustive should serve as a useful guide on the classifications of the various traits of existing MA research.

Classes	Characteristics	Example systems
1 st Generation	Global Search Paired with Local Search	<ul style="list-style-type: none"> i) A Canonical MA (Moscato 1989) ii) Adaptive global/local search (Hart 1994) iii) MA for Combinatorial Optimization (Land 1998) iv) Handling Computationally Expensive Problems (Ong, Nair et al. 2003) v) Multiobjective Permutation Flowshop Scheduling (Ishibuchi, Yoshida et al. 2003) vi) Fitness Landscape Analysis of MA (Merz 2004) vii) Robust Aerodynamic Design (Ong, Nair et al. 2006) viii) Evolutionary Gradient search (Arnold 2007) ix) Large-Scale Quadratic Assignment Problem (Tang, Lim et al. 2007) x) Evolutionary Lin-Kernighan for Traveling Salesman Problem (Meuth and Wunsch 2008) xi) Dynamic Optimization Problem (Wang, Wang et al. 2009) and many others.

2 nd Generation	Global Search with Multiple Local Optimizers. Memetic Information (Choice of Optimizer) Passed to Offspring. (Lamarckian Evolution)	<ul style="list-style-type: none"> i) Nurse Rostering Problem (Burke, Cowling et al. 2001) ii) Hyper-heuristic MA (Kendall, Soubeiga et al. 2002) iii) Structure prediction and structure comparison of proteins (Krasnogor, Blackburne et al. 2002) iv) Meta-Lamarckian MA (Ong and Keane 2004) v) Multimeme MA (Krasnogor and Smith 2005) vi) Adaptive Multi-Meme MA (Ong, Lim et al. 2006) vii) Multimeme Algorithm for Designing HIV Multidrug Therapies (Neri, Toivanen et al. 2007) viii) Agent-based Memetic Algorithm (Ullah, Sarker et al. 2007; Hasan, Sarker et al. 2008) ix) Diffusion Memetic Algorithm (Nguyen, Ong et al. 2008) and several others.
3 rd Generation	Global Search with Multiple Local Optimizers. Memetic Information (Choice of local Optimizer) Passed to Offspring. (Lamarckian Evolution). A mapping between evolutionary trajectory and choice of local optimizer is learned.	<ul style="list-style-type: none"> i) Co-evolution MA (Smith 2007) ii) Self-generation MA (Krasnogor and Gustafson 2004)
4 th Generation	Mechanisms of Recognition, Generalization, Optimization, and Memory are utilized.	Unknown

Table 1. Generational descriptions of memetic algorithms.

The mammalian brain exhibits hierarchical self-similarity, where neurons, groups of neurons, regions of the brain, and even whole lobes of the brain are connected laterally and hierarchically. Biological neurons are particularly well suited to this architecture; a single neuron serves as both a selection and learning mechanism. A neuron only fires when it receives significant input from one or more sources, and thus serves as a correlation detector. Additionally, it learns by modifying the weights of its inputs based on local information from firing rate, as well as global information from the chemical environment. Thus neurons activate when they encounter patterns that have made them fire before, and are able to adapt in delayed-reward situations due to global signals.

In laterally connected architectures, neuron groups can provide the function of clustering, as active neurons suppress the activity of their neighbors to pass their information down the processing chain, providing both selection and routing of information. The effect of this selectivity is that biological neural architectures route a spreading front of activation to different down-stream networks based on the similarity of the features present in the pattern of activation to previously presented patterns. As the activation front passes each neuron, the synaptic weights are changed based on local information – the firing rate of the neuron, the chemical environment, and the features present in the signal that activated the neuron, slightly changing how an individual neuron will respond at the next presentation of patterns (Beinenstock, Cooper et al. 1982).

Connected in loops, neurons provide short-term memory, process control and create temporally-delayed clustering. Combining loops and lateral connections at several levels of neuron groups (groups of neurons, groups of groups, etc) the neural architecture is able to exhibit increasing levels of selection, memory, and control. This is exactly the architecture that we see in the human cortex – a single cortical column contains recursion and lateral inhibition, and these cortical columns are arranged in a similar way, progressing in a fractal learning architecture up to the level of lobes, where sections of the brain are physically separated (Johansson

and Lansner 2007). This fractal architecture is similar to the N th-order meta-learning architecture described later in Section 4.

The brain inspired meta-learning memetic computational system is thus regarded here as a 4th generation memetic computational system. The novelty of the proposed meta-learning memetic system is highlighted below.

- i. In contrast to the 2nd generation memetic algorithms, there is no need to pre-define a pool of memes that will be used to refine the search. Instead memes are learned automatically - they are generalized information that passed between problem instances.
- ii. Since it satisfies all the three basic principles of an evolving system, it also qualifies as a 3rd generation memetic computational system. Unlike simple rule-based representation of meme used in co-evolution and self-generation MAs, the brain inspired meta-learning memetic computational system models the human brain that encodes each meme as hierarchies of cortical neurons (Johansson and Lansner 2007). With a self-organizing cortical architecture, meaningful information from recurring real-world patterns can be captured automatically and expressed in hierarchical nested relationships. A human brain stimulated by the recurrence of patterns, builds bidirectional hierarchical structures upward. The structure starts from the sensory neurons, through levels of cortical nodes and back down towards muscle activating neurons.
- iii. There exists a memory component to store the system's generalized patterns or structures of previously encountered problems - these elements could be thought of as memes.
- iv. Selection mechanisms are provided to perform association between problem features and previously generalized patterns that are likely to yield high-quality results.
- v. Meta-learning about the characteristics of the problem is introduced to construct meta-memes which are stored in the selection mechanism, allowing higher-order learning to occur automatically.
- vi. Memes and meta-memes in computing are conceptualized for higher-order learning as opposed to the typical definition of local search method used in all the works in memetic algorithm.

3. Schema-Meme Relationship

A genetic algorithm learns by passing schema (the genetic information of individuals) from generation to generation. Through natural selection and reproduction, useful schemata proliferate and are refined through genetic operators. The central concept of learning is that of the schema – a unit of information that is developed through a learning process (Holland 1975; Rumelhart 1980; Poli 2001). The typical ‘memetic algorithm’ uses an additional mechanism to modify schemata during an individual’s ‘lifetime,’ taken as the period of evaluation from the point of view of a genetic algorithm, and that refinement is able to be passed on to an individual’s descendants. The concept of schemata being passable just as behaviors or thoughts are passed on is what we term as memes – a meme being a unit of cultural information (Topchy and Punsch 2001; Ong and Keane 2004; Smart and Zhang 2004; Ong, Lim et al. 2006). Thus, memes can be thought of as an extension of schemata – schemata that are modified and passed on over a learning process. However, this distinction is a matter of scale. In a learning method, the current content of the representation could be called a schema, but when that information is passed between methods, it is more appropriately regarded as a meme.

This is analogous to the sociological definition of a meme (Dawkins 1989). In this form, a meme may contain certain food preparation practices, or how to build a home or which side of the road to drive on. Within the individuals of a generation, they are relatively fixed, but they are the result of a great deal of optimization, capturing the adaptations resulting from the history of a society. These cultural memes are passed from generation to generation of the population, being slightly refined at each step – new ingredients are added to the cooking methods, new building materials influence construction, traffic rules change, etc. The mechanism that allows this transformation is that of generalization (Rosca 1995; O'Neill and Ryan 1999; O'Neill and Ryan 2001). To communicate an internal schema from one individual to another, it must be generalized into a common representation – that of language in the case of human society. The specifics of the schema are of no great importance, as they would mean very little to an individual other than the originator due to the inherent differences between individuals. For instance, a description of the precise movements necessary to create a salad, such as the technique used to slice tomatoes and wash lettuce, is less important than the ingredients and general process of preparing the salad. Thus the salad recipe is a meme, a generalized representation of the salad, but the recipe alone is insufficient to produce the salad. The salad recipe is expressed only when it is put through the process of preparation, of acquiring and preparing the individual ingredients, and combining them according to the salad meme.

A meme may be thought of as generalized schema. Schemata are refined for an instance; memes are generalized to the extent of being transmissible between problem instances. To resolve the potential confusion that may arise, we put forth a loose definition of the term “Memetic Computation” – a paradigm of computational problem-solving that encompasses the construction of a comprehensive set of memes thus extending the capability of an optimizer to quickly derive a solution to a specific problem by refining existing general solutions, rather than needing to rediscover solutions in every situation.

4. A Framework for Higher Order Learning

A meta-learning system should be composed of 4 primary components – an optimizer, a memory, a selection mechanism, and a generalization mechanism, shown in Figure 2. The selection mechanism takes the features of a given problem as input, and performs a mapping to solutions in the memory that have an expected high quality. The memory stores previous or generalized solutions encountered by the system, and passes selected solution(s) on to the optimizer. The optimizer performs specialization and modification of solutions to optimize a given specific problem instance, while the generalization mechanism compares the resultant solution with existing solutions in memory, and either adds a new solution or modifies an existing solution. In memetic computation terms, the optimizer generates or modifies memes into schema, and then the generalization mechanism converts the schema back into memes for storage in memory. The selection mechanism provides a mapping about memes, providing recognition from a problem specification to a likely useful general solution, effectively utilizing internally represented meta-memes.

With these components, the architecture should be capable of exploiting information gained in previous problem sessions towards the solution of problems of increasing complexity. Integrating a cross-instance memory and a selection mechanism with an optimization method allows the recognition of a situation and the selection of previously utilized schema as likely high quality solution candidates. The optimization process then combines and refines these solution candidates to provide a good solution much faster than if the method had only random initial solutions. Once the solution is deployed, the selection method is trained to associate the situation (stimulus) with the solution (behavior) utilizing the fitness (reward) of the solution.

The process described above is itself a learning process, and thus could be augmented with increasingly higher level memory and selection methods, to allow complex, high-order solutions to be found. A sort of fractal meta-learning architecture of this type may be capable of virtually unlimited problem-solving capacity across a wide variety of problem domains.

The sequence of learning sessions matters greatly to the expression of complex behavior. By starting with simple problem instances and presenting successively more complex scenarios, the problem is decomposed, allowing solutions from sub-problems to be exploited, increasing the likelihood that higher level solutions will occur. Additionally, by training these simple solution components, a wider variety of high-level solutions can be trained more rapidly. For example, when training a dog, teaching him to ‘sit’ decreases the amount of training necessary for both ‘stay’ and ‘beg’ behaviors. This is analogous to the automatic construction of a ‘Society of Mind’ as described by (Minsky 1986).

When constructing optimization architectures, an issue of particular relevance is that of representation – how the schemata are stored. In population based algorithms schemata are stored as parameter strings, in neural networks, schemata are implicitly represented as interconnection weights, clustering methods store templates for categories, etc. How these schemata are expressed (and thereby their meaning) is dependent on the expression structure. In genetic algorithms a string is decoded into a trial problem solution, the weights in neural networks are utilized through weighted summation and passing through a transfer function. This division of representation prevents the simple utilization of schema across solution methods. To get disparate methods to work together, great care must be taken to modify both methods to utilize the same schema, which has been the subject of a great deal of research (Koza 1991; Angeline 1993; Merz and Freisleben 1997; Nguyen, Yoshihara et al. 2000; Abramson and Wechsler 2001; Baraglia, Hidalgo et al. 2001; Mulder and Wunsch 2003; Ong and Keane 2004; Agarwal, Lim et al. 2005; Dang and Zhang 2005).

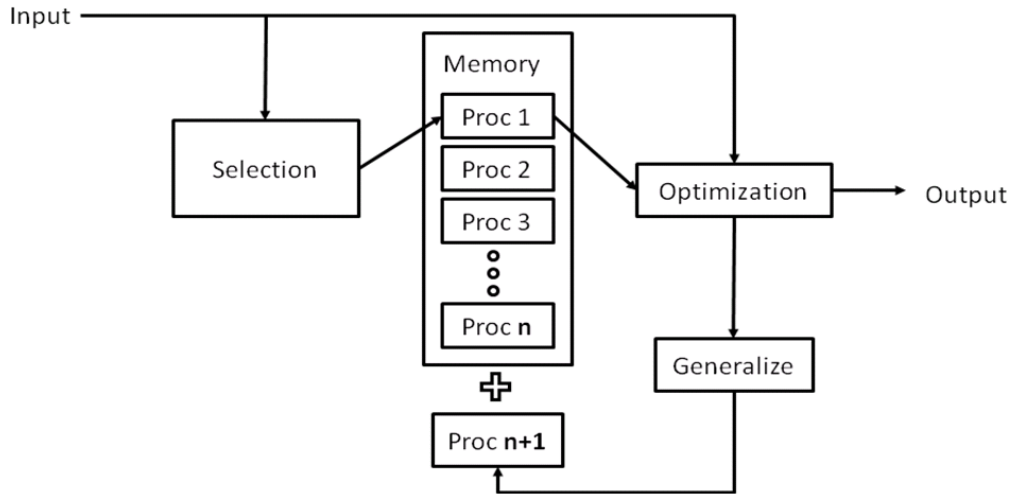


Figure 2. Meta-Learning Architecture

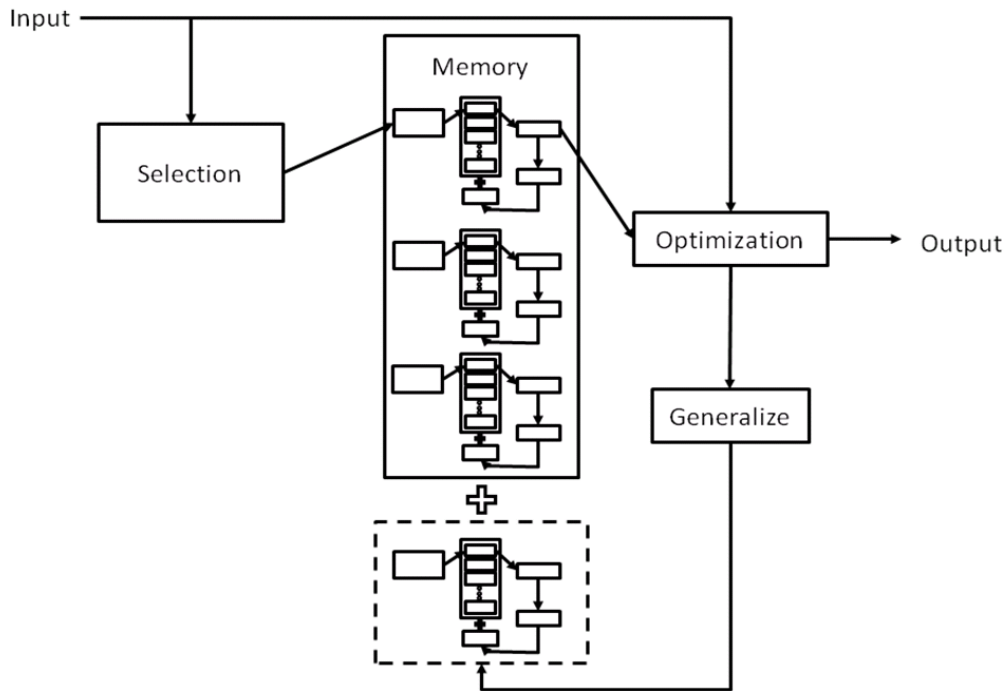


Figure 3. Meta-Meta Learning

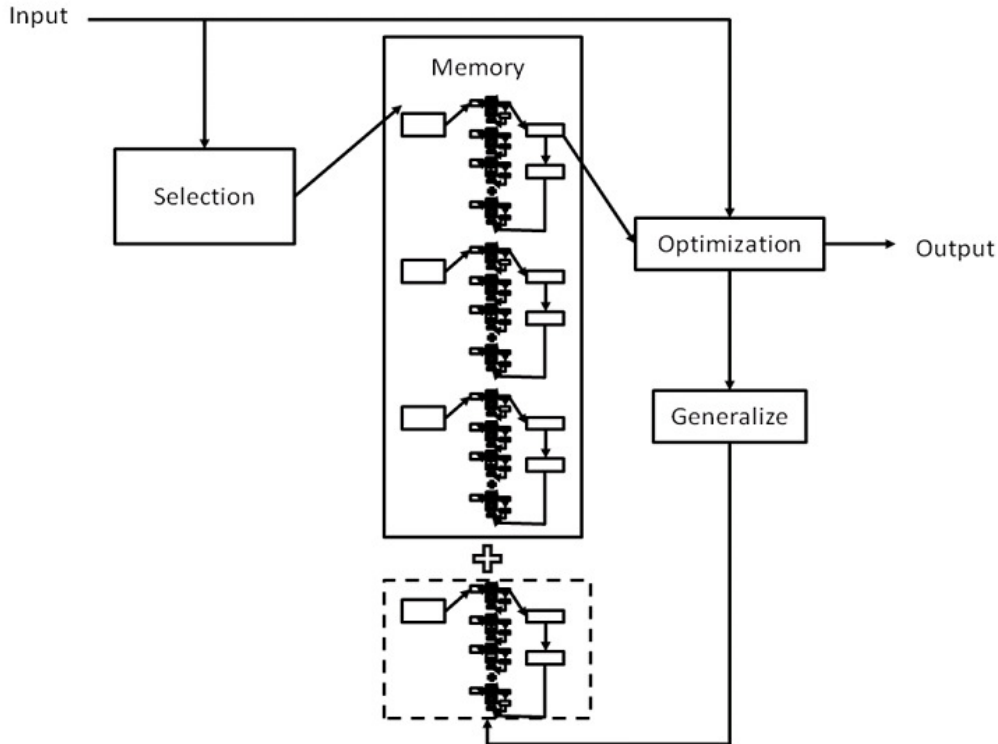


Figure 4. N^{th} -order Meta Learning

First order learning methods consist of a single algorithm that modifies schema to optimize a system. Individually, all classical machine learning methods fall into this category. Meta-learning or second-order methods learn about the process of learning, and modify the learning method, which in turn modifies schema. A simple illustration of a meta-learning architecture is presented in Figure 2. In this figure, schemata are represented as “procedures”, which are stored in memory. A problem is presented to the architecture, and a selection mechanism chooses likely valuable schema from memory, which are then modified to the particular problem instance. High-value schema are then generalized and saved back into memory, and the selection mechanism then learns an association between characteristics of the problem instance and schema that yielded positive results.

These second order methods should be able to be combined with other methods or layers to produce third-order methods and so on to order N , as illustrated in Figures 3 and 4. To produce higher order methods, information gained in one problem instance should be utilized to provide a partial solution to another similar problem instance allowing the system as a whole to take advantage of previous learning episodes.

5. Even-Parity Example

To demonstrate the principles and advantages of meta-learning, we examine its application to the even and odd parity problems, standard benchmarks for genetic programming and automatic function definition methods (Koza 1992). We propose a hypothetical genetic programming system utilizing a set of Boolean operators to construct individuals implementing the even or odd parity functions (XOR and XNOR, respectively). We analyze two cases of the evolution of the three-input XOR function, both starting with populations implementing the two-input XOR function, with and without the abstraction that is inherent in a meta-learning system. A third case is presented illustrating the functionality of a simple selection mechanism on the odd-parity function.

a) Problem Overview

Koza described the even parity problem below.

“The Boolean even-parity function of k Boolean arguments returns \top (True) if an odd number of its arguments are \top , and otherwise returns \perp (False). The concatenation of this returned bit to the original string making the total string even, hence even-parity.

In applying genetic programming to the even-parity function of k arguments, the terminal set T consists of the k Boolean arguments D_0, D_1, D_2, \dots involved in the problem, so that

$$T = \{D_0, D_1, D_2, \dots\}.$$

The function set F for all the examples herein consists of the following computationally complete set of four two-argument primitive Boolean functions:

$$F = \{\text{AND, OR, NAND, NOR, NOT}\}.$$

The Boolean even-parity functions appear to be the most difficult Boolean functions to find via a blind random generative search of expressions using the above function set F and the terminal set T . For example, even though there are only 256 different Boolean functions with three arguments and one output, the Boolean even-3-parity function is so difficult to find via a blind random generative search that we did not encounter it at all after randomly generating 10,000,000 expressions using this function set F and terminal set T . In addition, the even-parity function appears to be the most difficult to learn using genetic programming using this function set F and terminal set T (Koza 1992).”

The odd-parity function is similarly constructed, returning true if an even number of its arguments are true, and otherwise returning false.

In genetic programming (GP), the genome of an individual is represented as a tree structure, where operations are applied at branches, and the leaves are constants and problem parameters. An illustration of a functional represented as tree structure is shown in Figure 5 (Koza 1989; Koza 1992). One advantage of GP is that the results can be easily human interpretable and formally verifiable, a quality that is not present in many other computational intelligence methods (Rosca 1995).




Figure 5. Illustration of function representation as tree structure.

The even-2-parity function is simply the XOR function, which is itself a composition of the terminal set functions in one simple possible configuration:

$$a \text{ XOR } b = (a \text{ OR } b) \text{ AND } (a \text{ NAND } b)$$

Using a tree representation, the XOR function is shown in Figure 6.




Figure 6. XOR tree representation.

Constructing the even-3-parity function using only these primitives is more difficult, but follows a similar pattern, illustrated below and in Figure 7:

$$\text{XOR}(a, b, c) = (((a \text{ OR } b) \text{ AND } (a \text{ NAND } b)) \text{ OR } c) \text{ AND } (((a \text{ OR } b) \text{ AND } (a \text{ NAND } b)) \text{ NAND } c)$$




Figure 7. Three-input XOR tree representation.

Note that the three-input XOR structure relies on the recursive use of the two-input XOR function, replacing the 'a' nodes with XOR nodes, and re-assigning the top-level 'b' nodes to be the 'c' variable.




Figure 8. Simplified two-input XOR.

Note that if a 2-bit XOR function is defined explicitly as in Figure 8, the even-3-parity function becomes greatly simplified, as written below and shown in Figure 9.

$$\text{XOR}(a, b, c) = (a \text{ XOR } b) \text{ XOR } c$$

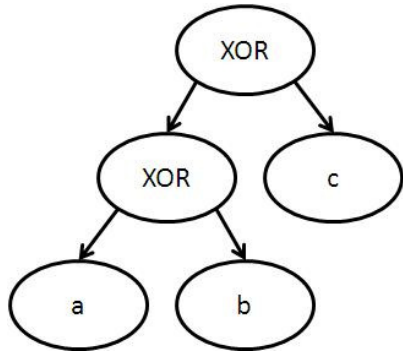


Figure 9. Simplified three-input XOR.

b) Case 1 – Non-Meta XOR3 Evolution

Taking a genetic programming system as an example, in a non-meta learning system, evolution of the XOR3 function must proceed through at least two generations. To further expand on our illustration, we consider the best case scenario whereby all the individuals in the population incorporate the simplified XOR function, as shown in Figure 10.

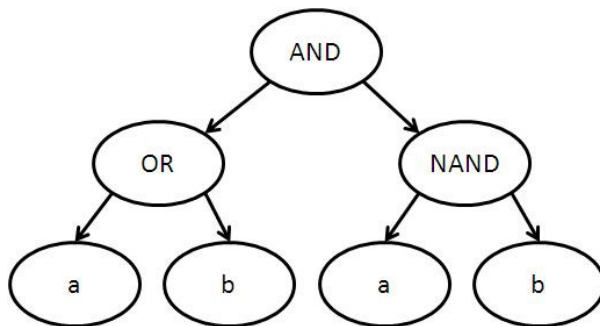


Figure 10. Initial non-meta learning XOR2 individual.

As there are 4 leaf nodes out of 7 total nodes, the probability of selecting a leaf node for crossover (P_{L1}) is $4/7$. Assuming a uniform population of individuals implementing XOR2 (translating to a 100% probability of choosing another XOR2 individual for crossover) the probability of selecting the root node of another individual to replace the selected leaf node is (P_{F1}) $1/7$.




Figure 11. Intermediate step in development of 3-bit XOR function after a single generation.

Then, the evolutionary process must select one of the two top-level 'b' nodes for mutation from the tree which has a total of thirteen nodes, thus the probability of selecting one correct leaf for mutation (P_{M1}) is $2/13$. Choosing from the eight possible node types (the combination of terminal set and functional set), the probability of selecting the correct 'c' variable (P_{V1}) is $1/8$.

At this point the evolutionary reproduction steps are completed, and the individual shown in Figure 11 is evaluated. This partial XOR3 function is not yet complete, but it correctly completes one test case more than the XOR2 function, which may give it an evolutionary advantage. Assuming that the individual survives to the next generation and is again selected as a parent with 100% probability, an additional reproduction step must be completed to yield an XOR3 function.

Now the correct leaf node must be selected for crossover, but this time there is only one node, the 'a' node at a depth of three, from the thirteen possible nodes, so the probability of selecting the correct leaf node for crossover (P_{L2}) is $1/13$. Once again, assuming all other individuals in the population still implement the XOR2 function in Figure 8, the probability of selecting the root of another XOR2 individual to replace the leaf (P_{F2}) is $1/7$. At the completion of crossover, the total number of nodes in the tree becomes eighteen. At the mutation step, the remaining 'b' node at depth three must be selected, and the probability of selecting correct leaf for mutation (P_{M2}) is $1/18$. Completing the XOR3, the probability of selecting the correct variable from the total set of node types (P_{V2}) is $1/8$. The completed three-input XOR function is illustrated earlier in Figure 9.

Ignoring changes in the population and evolutionary survivability, the probability of transitioning from XOR2 to XOR3 in two generations without meta-learning is calculated below.

$$P_{\text{xor3_nonmeta}} = P_{L1} * P_{F1} * P_{M1} * P_{V1} * P_{L2} * P_{F2} * P_{M2} * P_{V2} = 1.19 \times 10^{-7}$$

where,

- P_{L1} is the probability of a leaf node selection for crossover during the first generation,
- P_{F1} is the probability of functional root selection for crossover during the first generation,
- P_{M1} is the probability of proper leaf selection for mutation during the first generation,
- P_{V1} is the probability of proper variable selection for mutation during the first generation,
- P_{L2} is the probability of a leaf node selection for crossover during the second generation,
- P_{F2} is the probability of functional root selection for crossover during the second generation,
- P_{M2} is the probability of proper leaf selection for mutation during the second generation,
- P_{V2} is the probability of proper variable selection for mutation during the second generation.

Note that this ignores the significant influence of relative fitness, generational selection, parent selection, probability of application of crossover/mutation operators and population influence and may be interpreted as a kind of upper-bound on the probability that a two-input XOR individual will develop into a three-input XOR without the abstraction capability of meta-learning.

c) Case 2 – Meta-Learning XOR3 Evolution

In this case we assume a meta-learning system that has already learned a two-input XOR function, performed generalization and added this to the function set ($F = \text{AND, OR, NAND, NOR, NOT, XOR2}$). The probability that the system will transition from XOR2 to XOR3 is calculated using only the mutation step.

With a population uniformly initialized with the two-input XOR and an individual selected from this population, illustrated in Figure 8, the probability of selecting a leaf node for mutation (P_L) is $2/3$ as the simplified XOR tree has only 3 nodes, and two of them are terminals. Having selected a terminal, the probability of selecting the XOR2 function from the node set of six functions and three terminals to replace the leaf node (P_F) is $1/9$. Assuming a recursive mutation process, two new leaf nodes must be selected, and they must contain variables not yet used by the tree to produce a three-input XOR. The probability of selecting the correct terminal node is $1/9$, and this process must be repeated twice, so the probability of selecting two correct terminal nodes (P_V) is $(1/9)^2$ or $1/81$. Using only one generation the three-input XOR can be developed in a meta-learning system.

$$\text{Probability of XOR3 from XOR2: } P_{\text{xor3_meta}} = P_L * P_F * P_V = 0.000914$$

where,

- P_L is the probability of a leaf node selection for mutation,

P_F is the probability of XOR2 function selection for mutation,

P_V is the probability of proper leaf selection for mutation.

Note that using meta-learning, the three-input XOR can also occur with a crossover and a mutation, where the non-meta learning system must utilize two full generations. Also note that though the size of the functional set has increased, the number of changes necessary to place an upper-bound on the probability of a three-input XOR occurring has been substantially decreased, allowing the evolutionary process to focus on high-level changes.

Thus in a large population, the XOR3 function may occur in a single generation with a meta-learning system, where a non-meta learning system must take at least two generation and probably many thousands of evaluations to evolve an XOR3.

d) Case 3 – Selection and Odd-Parity Evolution

To demonstrate the advantages of the complete meta-learning procedure, we first present the 2-bit even-parity problem to a theoretical meta-learning system, then the 2-bit odd-parity problem, and finally the 3-bit even-parity problem. The selection mechanism shall have 2 inputs – the first is activated only when the system is operating on the even-parity problem, the second is activated only when operating on the odd-parity problem. Initially, the memory is empty, so the optimizer is initialized with random solutions.

Presented with the even-2-parity problem, the optimizer outputs a resulting solution that performs the XOR function – “D0 XOR D1”, where D0 and D1 are the Boolean arguments of the input. This function is passed to the generalization mechanism, which removes the absolute references to the Boolean arguments, replacing them with dummy variables ‘A’ and ‘B’, resulting in the function “A XOR B”. This generalized XOR function is then added to the memory, making the function available as a primitive. The functional set becomes:

$F = \{AND, OR, NAND, NOR, NOT, XOR\}$.

The selection mechanism is updated to learn an association between the active ‘even-parity’ input and the new memory element. At this point the procedure and difference in optimization would be no different than if the optimizer were operating without the rest of the meta-learning architecture.

Next, the odd-2-parity problem is presented, the ‘odd-parity’ input is activated on the selector mechanism, and having no other elements to select, the sole item in memory (the generalized “A XOR B” function) is selected to initialize the state of the optimizer. The optimizer replaces the dummy variables with references to the Boolean arguments and begins optimization. As only a small modification is necessary, the addition of the NOT primitive function at a high-level to create an XNOR function, the optimizer has a high probability of quickly finding a perfect solution to the odd-2-parity problem. This differs from a randomly initialized optimizer as there would be a lower probability of finding a good solution due to the need to explore more modifications. Once the meta-learning optimizer finds the solution, the generalization, memory insert, and selection training steps are repeated for the XNOR function:

$F = \{AND, OR, NAND, NOR, NOT, XOR, XNOR\}$.

Finally, the even-3-parity problem is presented to the meta-learning architecture. The selection ‘even-parity’ input is activated, and the associated XOR memory element is used to initialize the optimizer state. The optimizer replaces the XOR dummy variables with argument references, and begins the optimization process. The optimizer need only make the relatively small change of cascading the XOR function to produce a 3-input XOR function, where a raw optimization function without a memory or selection method would need to evaluate and modify many combinations of the original 5 functional primitives to arrive at a good solution. Thus the meta-learning architecture should be able to arrive at high-value solutions rapidly by exploiting previously generated solution to construct high-level solutions.

In this example the memory component stores generalized solutions to previously encountered problems - these elements could be thought of as memes, as they are solutions that are passed between problem instances. The selection mechanism performs association between problem features and solutions that are likely to yield high-value results. By not only providing the input data to the problem, but additional meta-data about the characteristics of the problem, the meta-learning architecture can construct meta-memes which are stored in the selection mechanism, allowing higher-order learning to occur automatically.

6. Traveling Salesman Problem

The Traveling Salesman Problem (TSP) is a standard combinatorial optimization problem used for the design and evaluation of optimization methods (Lin and Kernighan 1973; Wang, Maciejewski et al. 1998; Nguyen, Yoshihara et al. 2000; Baraglia, Hidalgo et al. 2001; Tsai, Yang et al. 2002; Applegate, Cook et al. 2003; Mulder and Wunsch 2003; Wunsch and Mulder 2003; Tsai, Yang et al. 2004; Dang and Zhang 2005; Xu and Wunsch 2005; Agarwal, Lim et al. 2007; Nguyen, Yoshihara et al. 2007). TSP optimization algorithms have a wide range of applications including job scheduling, DNA sequencing, traffic management, and robotic path planning. To further illustrate the capabilities of the meta-learning design paradigm, an example is presented using instances of the TSP.

To apply meta-learning to the TSP problem, the schema of the problem must be identified. Here the schema takes the form of the ordering of points in a tour. The addition of a clustering method to divide and conquer the TSP has been shown to greatly accelerate the solution of the TSP (Meuth and Wunsch 2008). With this addition, the overall schema for the optimizer consists of the combination of cluster templates, tour point ordering, and the locations of points. This schema must be generalized to create a meme, which is trivial for the cluster templates, but more challenging for the tour ordering and points' locations. The problem is further complicated by the necessity to generalize tours to be applicable over multiple scales.

For this application, a meme consists of a small ordered tour, containing small, limited number of points. To generalize the meme, the centroid of the group is calculated and subtracted from each point, making the centroid the origin of the group. The coordinates of each point are then normalized by distance from the origin. This projects the points into unit-space, and allows comparisons across multiple scales. Each TSP-meme serves as a pre-optimized tour template. Each point in the TSP-meme can represent a real point in the problem instance, or the centroid of a group of points, itself represented by a meme.

Given an instance of the TSP, the meta-TSP algorithm utilizes a clustering method to divide the problem into sub-problems, and divides those sub-problems into sub-sub problems and so on, until a threshold for sub-problem size is reached. The relationships between sub-problems are recorded in a tree-representation. Each of these sub-problems is generalized, and compared against the recorded memes for existing solutions.

The recognition mechanism must be able to detect structurally similar sub-problems. For this experiment the matching mechanism compares two normalized sub-problems by finding the nearest corresponding points between the memes, and calculating the mean squared error between these points.

If a match is found in memory, the existing meme-solution (a point ordering) is copied to the current sub-problem, and the sub-problem updates the meme by refining template point positions. If no match exists in memory, the sub-problem is solved as accurately as possible. With a small enough problem threshold, exact solutions to sub-problems can be found, depending on computational resources available. The sub-problem is then stored in memory as a new meme. After all the sub-problems are solved, they are combined into a global tour by collapsing the problem-tree, and utilizing a simple $O(n)$ merge algorithm as detailed in (Mulder and Wunsch 2003).

To illustrate this process, an example is given utilizing a simple instance of the TSP, shown in Figure 12. A first pass of clustering is shown in Figure 13. Note that cluster M3 contains many points, and that a single point has been left out of the clusters for illustrative purposes. A second pass further divides cluster M3 into clusters M5, M6, and M7, as shown in Figure 14. The final clustering pass assigns all clusters to a global cluster, M8, in Figure 15. The hierarchy of clusters, and thereby sub-problems, is denoted by the cluster tree in Figure 16.

At this stage, each sub-problem is optimized independently, as shown in Figure 17. Note that some of the sub-problems contain references to other sub-problems, particularly M3 and M8. The centroids of sub-problems are utilized for optimization and solution, representing sub-problems as a whole. During the course of optimization, each sub-problem is normalized, and compared with previously computed, normalized solutions in the memory. These memes can be stored across instances, building a large library of pre-computed solutions that can be deployed to yield high quality solutions rapidly. Sub-problems of a global problem instance can be thought of as new problem instances, and pre-computed solutions that are generated during the calculation of a global instance can be applied across sub-problems.

For example, the normalized versions of M2 and M4 would be very similar in structure, and once M2 is computed, the structural similarity of the sub-problems would be recognized, and the ordering of points for M4

need not to be computed, only copied from M2 to M4. The same process applies across scales and global problem instances.




Figure 12. Small TSP instance of approximately 30 points.




Figure 13. TSP Instance after first clustering pass. Each cluster initializes a meme, labeled with "M#" and a "+" denoting the centroid.




Figure 14. Second clustering pass. Note the new clusters, M5, M6, and M7.

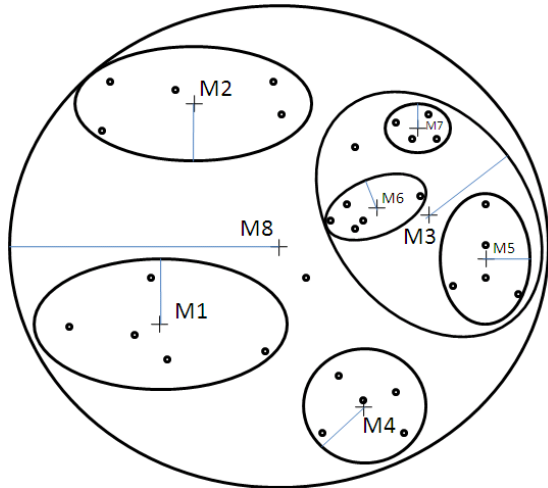


Figure 15. Final clustering pass, with global cluster M8




Figure 16. Tree of sub-problems (clusters)

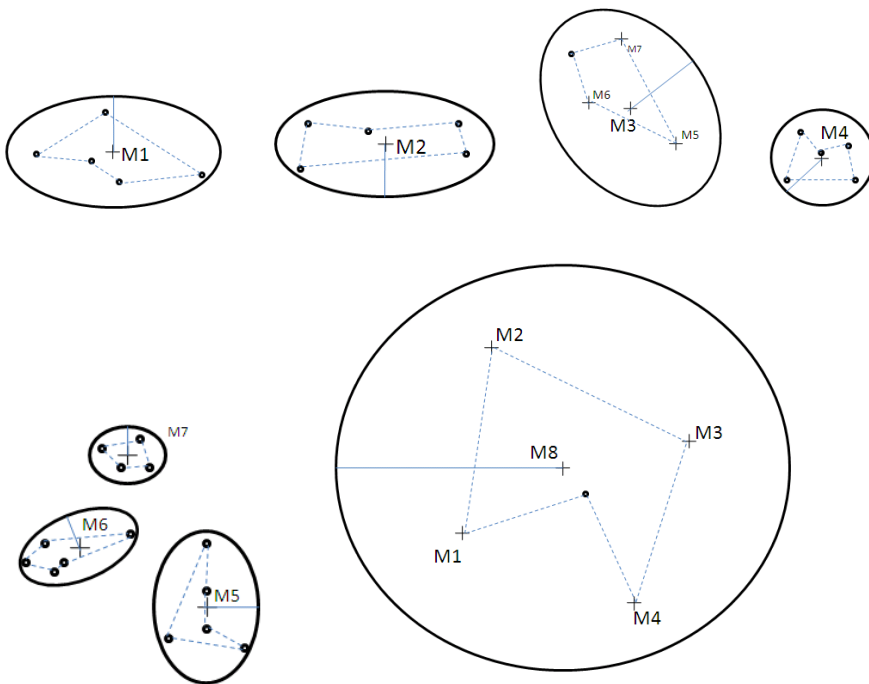


Figure 17. Completed memes, M1 through M8. Super-clusters reference the centroids of sub-clusters. Note that memes M2 and M4 are similar in structure, but not scale.




Figure 18. Memes M5, M6, and M7 are merged into M3.




Figure 19. Memes M1, M2, M3, M4 are merged into M8




Figure 20. Completed tour.

When all sub-problems are completed, the problem hierarchy is collapsed by de-referencing sub-problems and incrementally merging them with higher level tours. This is accomplished by choosing the closest set of two vertices in the sub-problem to any two vertices in the higher level tour. To avoid an $O(n^2)$ operation, a small neighborhood of vertices from the super-tour is chosen based on proximity to the centroid of the sub-tour. This neighborhood of super-tour vertices is compared to each vertex in the sub-tour to find the best match. A result of this merge operation is illustrated in Figures 18 and 19. Figure 19 shows the final merge of all complete sub-tours into a final tour. The completed tour is shown in Figure 20.

The computational complexity of the proposed method is expected to be very efficient at $O(n \log(n))$ improving with linearly decreasing complexity as the library of pre-optimized solutions grows, decreasing the amount of optimization to be performed on a given TSP instance.

The qualitative performance of this method is the subject of future development. The algorithm presented here serves as an example of meta-learning driven design, incorporating mechanisms of memory, selection, optimization, and generalization.

7. Conclusion

The desire for a new and robust computational intelligence paradigm spans many problem domains, including real time robotic systems which must deal with increasing complexity on a daily basis, deep data mining such as natural language processing with applications in information retrieval and machine understanding, human-computer interaction, and long-term optimization. These new, complex frontiers of machine learning and optimization could all benefit from the higher order memetic computing methods described here.

We have presented an overview of important definitions and architectures in memetic computing and have attempted to illustrate the power of next-generation memetic algorithms. The primary difficulty of designing meta-learning systems lies in the construction of valid scale-invariant representations which enable the simple construction of selection, generalization, and memory mechanisms. By providing generalization, memory, optimization, and selection mechanisms, a meta-learning architecture can operate on high-level features of a problem instance, selecting generalized solutions that have been used previously with high utility in the problem context. Utilizing these features, a system should be able to learn not only the solution to a problem, but learn *about* solving problems. Such a system may enable a quantum leap in the performance of real-world adaptive systems as they provide the central components of meta-adaptive systems to be constructed.

8. References

- Abramson, M. and H. Wechsler (2001). Competitive reinforcement learning for combinatorial problems. International Joint Conference on Neural Networks, 2001. Proceedings. IJCNN '01. .
- Agarwal, A., M.-H. Lim, M.-J. Er and C.-Y. Chew (2005). ACO for a new TSP in region coverage. IEEE/RSJ International Conference on Intelligent Robots and Systems.
- Agarwal, A., M.-H. Lim, M. J. Er and T. N. Nguyen (2007). "Rectilinear workspace partitioning for parallel coverage using multiple UAVs." Advanced Robotics **21**(1).
- Angeline, P. J. (1993). Evolutionary Algorithms and Emergent Intelligence. Columbus, OH, Ohio State University. **Doctoral Thesis**.
- Applegate, D., W. Cook and A. Rohe (2003). "Chained Lin-Kernighan for Large Traveling Salesman Problems." INFORMS Journal on Computing **15**(1): 82-92.
- Arnold, D. V. a. S., R. (2007). "Evolutionary Gradient Search Revisited." IEEE Transactions on Evolutionary Computation **11**(4): 480 - 495.
- Baraglia, R., J. I. Hidalgo and R. Perego (2001). "A hybrid heuristic for the traveling salesman problem." IEEE Transactions on Evolutionary Computation **5**(6): 613-622.
- Beinestock, E. L., L. Cooper and P. Munro (1982). "Theory for the development of neuron selectivity: orientation specificity and binocular interaction in the visual cortex." Journal of Neuroscience **2**(1): 32-48.
- Burke, E., P. Cowling, P. D. Causmaecker and G. Berghe (2001). "A Memetic Approach to the Nurse Rostering Problem." Applied Intelligence **15**(3): 199-214.
- Dang, J. and Z. Zhang (2005). A Polynomial Time Evolutionary Algorithm for the Traveling Salesman Problem. International Conference on Neural Networks and Brain.
- Dawkins, R. (1989). The Selfish Gene, Oxford University Press, USA.
- Francois, O. and C. Lavergne (2001). "Design of evolutionary algorithms-A statistical perspective." Evolutionary Computation, IEEE Transactions on **5**(2): 129-148.
- Gaudiot, J.-L., J.-Y. Kang and W. W. Ro (2005). "Techniques to Improve Performance Beyond Pipelining: Superpipelining, Superscalar, and VLIW, ." Advances in Computers(63).
- Hart, W. E. (1994). Adaptive Global Optimization with Local Search, University of California.
- Hasan, S. M. K., R. Sarker, D. Essam and D. Cornforth (2008). "Memetic Algorithms for Solving Job-Shop Scheduling Problems." Memetic Computing Journal.
- Holland, J. H. (1975). Adaptation in natural and artificial systems. Ann Arbor, University of Michigan Press.

- Ishibuchi, H., T. Yoshida and T. Murata (2003). "Balance between genetic search and local search in memetic algorithms for multiobjective permutation flowshop scheduling." IEEE Transactions on Evolutionary Computation **7**(2): 204–223.
- Johansson, C. and A. Lansner (2007). "Towards cortex sized artificial neural systems." Neural Networks **20**(1): 48-61.
- Kazarlis, S. A., S. E. Papadakis, J. B. Theocharis and V. Petridis (2001). "Microgenetic algorithms as generalized hill-climbing operators for GA optimization." Evolutionary Computation, IEEE Transactions on **5**(3): 204-217.
- Kendall, G., E. Soubeiga and P. Cowling (2002). Choice function and random hyperheuristics. 4th Asia-Pacific Conference on Simulated Evolution And Learning: 667–671.
- Kolodner, J. (1993). Case-Based Reasoning. San Francisco, CA, USA, Morgan Kaufmann Publishers Inc.
- Koza, J. R. (1989). Hierarchical genetic algorithms operating on populations of computer programs. International Joint Conference on Artificial Intelligence, Morgan Kaufman Publishers.
- Koza, J. R. (1991). Evolution and co-evolution of computer programs to control independent-acting agents. From Animals to Animats: Proceedings of the First International Conference on Simulation of Adaptive Behavior.
- Koza, J. R. (1992). The genetic programming paradigm: Genetically breeding populations of computer programs to solve problems. Dynamic, Genetic and Chaotic Programming, John Wiley: 201-321.
- Koza, J. R. (1992). Hierarchical Automatic Function Definition in Genetic Programming. Foundations of Genetic Algorithms 2, Morgan Kaufmann: 297-318.
- Krasnogor, N. (2002). Studies on the Theory and Design Space of Memetic Algorithms. Faculty Comput., Math. Eng. Bristol, U.K., Univ. West of England. **Ph.D.**
- Krasnogor, N., B. Blackburne, J. D. Hirst and E. K. Burke (2002). Multimeme algorithms for the structure prediction and structure comparison of proteins. Proc. Parallel Problem Solving From Nature, Springer. **Lecture Notes in Computer Science**.
- Krasnogor, N. and S. Gustafson (2004). "A Study on the use of "self-generation" in memetic algorithms." Natural Computing **3**(1): 53 - 76.
- Krasnogor, N. and J. Smith (2005). "A tutorial for competent memetic algorithms: model, taxonomy, and design issues." Evolutionary Computation, IEEE Transactions on **9**(5): 474-488.
- Kuncheva, L. I. and L. C. Jain (2000). "Designing classifier fusion systems by genetic algorithms." Evolutionary Computation, IEEE Transactions on **4**(4): 327-336.
- Land, M. W. S. (1998). Evolutionary Algorithms with Local Search for Combinatorial Optimization, UNIVERSITY OF CALIFORNIA.
- Lee, J. T., E. Lau and H. Yu-Chi (2001). "The Witsenhausen counterexample: a hierarchical search approach for nonconvex optimization problems." Automatic Control, IEEE Transactions on **46**(3): 382-397.
- Lenat, D. and R. V. Guha (1989). Building Large Knowledge-Based Systems, Addison-Wesley.
- Lin, S. and B. W. Kernighan (1973). "An Effective Heuristic Algorithm for the Traveling Salesman Problem." Operations Research **21**(2): 498-516.
- Merz, P. (2004). "Advanced Fitness Landscape Analysis and the Performance of Memetic Algorithms." Evolutionary Computation **12**(3): 303--325.
- Merz, P. and B. Freisleben (1997). Genetic Local Search for the TSP: new results. IEEE Conference on Evolutionary Computation.
- Meuth, R. J. and D. C. Wunsch, II (2008). Divide and Conquer Evolutionary Tsp Solution for Vehicle Path Planning. Congress on Evolutionary Computation (WCCI'08).
- Milano, M. and A. Roli (2004). "MAGMA: a multiagent architecture for metaheuristics." Systems, Man, and Cybernetics, Part B, IEEE Transactions on **34**(2): 925-941.
- Minsky, M. (1986). The Society of Mind, Simon & Schuster Inc.
- Moscato, P. (1989). On evolution, search, optimization, genetic algorithms and martial arts: Towards memetic algorithms, Caltech Concurrent Computation Program, C3P Report, 826.
- Mulder, S. and D. C. Wunsch (2003). "Million City Traveling Salesman Problem Solution by Divide and Conquer Clustering with Adaptive Resonance Neural Networks." Neural Networks.
- Neri, F., J. Toivanen, G. L. Cascella and Y. Ong (2007). "An Adaptive Multimeme Algorithm for Designing HIV Multidrug Therapies." IEEE ACM TRANSACTIONS ON COMPUTATIONAL BIOLOGY AND BIOINFORMATICS **4**(2): 264.
- Nguyen, H. D., I. Yoshihara, K. Yamamori and M. Yasunaga (2000). Modified Edge Recombination Operators of Genetic Algorithms for the Traveling Salesman Problem. 26th Annual Conference of the IEEE Industrial Electronics Society, 2000.
- Nguyen, H. D., I. Yoshihara, K. Yamamori and M. Yasunaga (2007). "Implementation of an Effective Hybrid GA for Large Scale Traveling Salesman Problems." IEEE Transactions on Systems, Man and Cybernetics Part B **37**(1): 92-99.

- Nguyen, Q.-H., Y.-S. Ong and M.-H. Lim (2008). Non-Genetic Transmission of Memes by Diffusion. 10th Annual Conference on Genetic and Evolutionary Computation (GECCO'08). Atlanta, GA.
- Nguyen, Q. H., Y. S. Ong and N. Krasnogor (2007). A study on the design issues of Memetic Algorithm IEEE Congress on Evolutionary Computation Singapore, IEEE: 2390-2397.
- Nguyen, Q. H., Y. S. Ong and M. H. Lim (2008). Non-genetic transmission of memes by diffusion.
- O'Neill, M. and C. Ryan (1999). Automatic Generation of High Level Functions using Evolutionary Algorithms. 1st International Workshop on Soft Computing Applied to Software Engineering, Limerick University Press.
- O'Neill, M. and C. Ryan (2001). "Grammatical Evolution." IEEE Transactions on Evolutionary Computation **5**(4): 349-358.
- Ong, Y.-S. and A. J. Keane (2004). "Meta-Lamarckian Learning in Memetic Algorithms." IEEE Transactions on Evolutionary Computation **8**(2).
- Ong, Y.-S., M.-H. Lim, N. Zhu and K.-W. Wong (2006). "Classification of Adaptive Memetic Algorithms: A Comparative Study." IEEE Transactions on Systems, Man and Cybernetics Part B **36**(1).
- Ong, Y. S. and A. J. Keane (2004). "Meta-Lamarckian learning in memetic algorithms." IEEE Transactions on Evolutionary Computation **8**(2): 99--110.
- Ong, Y. S., P. B. Nair and A. J. Keane (2003). "Evolutionary Optimization of Computationally Expensive Problems via Surrogate Modeling." AIAA Journal **41**(4): 687--696.
- Ong, Y. S., P. B. Nair and K. Y. Lum (2006). "Max-Min Surrogate-Assisted Evolutionary Algorithm for Robust Aerodynamic Design." IEEE Transactions on Evolutionary Computation **10**(4): 392--404.
- Poli, R. (2001). "Exact Schema Theory for Genetic Programming and Variable-Length Genetic Algorithms with One-Point Crossover." Genetic Programming and Evolvable Machines **2**(2): 123-163.
- Rosca, J. P. (1995). Genetic Programming Exploratory Power and the Discovery of Functions. Conference on Evolutionary Programming, MIT Press.
- Rumelhart, D. E. (1980). Schemata: The building blocks of cognition. . Theoretical Issues in Reading And Comprehension. B. B. R.J. Sprio, & W.F. Brewer. NJ, Erlbaum.
- Shahaf, D. and E. Amir (2007). "Towards a Theory of AI Completeness." 8th Interational Symposium on Logic Formalizations of Commonsense Reasoning.
- Smart, W. and M. Zhang (2004). Applying Online Gradient Descent Search to Genetic Programming for Object Recognition. Second workshop on Australasian information security, Data Mining and Web Intelligence, and Software Internationalisation, Dunedin, New Zealand.
- Smith, J. E. (2007). "Coevolving Memetic Algorithms: A Review and Progress Report." IEEE Transactions on Systems, Man, and Cybernetics, Part B **37**(1): 6-17.
- Tang, J., M. H. Lim and Y. S. Ong (2007). "Diversity-Adaptive Parallel Memetic Algorithm for Solving Large Scale Combinatorial Optimization Problems." Soft Computing **11**(9): 873-888.
- Topchy, A. and W. F. Punsch (2001). Faster Genetic Programming based on Local Gradient Search of Numeric Leaf Values. Genetic and Evolutionary Computation Conference.
- Tsai, H.-K., J.-M. Yang and C.-Y. Kao (2002). Solving traveling salesman problems by combining global and local search mechanisms. Conference on Evolutionary Computation.
- Tsai, H.-K., J.-M. Yang and C.-Y. Kao (2004). "An evolutionary algorithm for large traveling salesman problems." IEEE Transactions on Systems, Man and Cybernetics Part B **34**(4): 1718-1729.
- Ullah, A. S. S. M. B., R. A. Sarker, D. Cornforth and C. Lokan (2007). An agent-based memetic algorithm (AMA) for solving constrained optimization problems. IEEE Congress on Evolutionary Computation, Singapore.
- Wang, H., D. Wang and S. Yang (2009). "A memetic algorithm with adaptive hill climbing strategy for dynamic optimization problems." Soft Computing.
- Wang, L., A. A. Maciejewski, H. J. Seigel and V. P. Roychowdhury (1998). A comparative study of five parallel genetic algorithms using the traveling salesman problem. First Merged International Conference and Symposium on Parallel and Distributed Processing.
- Wills, L. M. and J. Kolodner (1994). "Towards More Creative Case-Based Design Systems." Proceedings of the Twelfth Annual National Conference on Artificial Intelligence (AAAI-94): 50-55.
- Woldpert, D. H. and W. G. Macready (1997). "No Free Lunch Theorms for Optimization." IEEE Transactions on Evolutionary Computation **1**(1): 67-82.
- Wunsch, D. C. and S. Mulder (2003). Using Adaptive Resonance Theory and Local optimization to divide and conquer large scale traveling salesman problems. International Joint Conference on Neural Networks.
- Xin, Y. (1999). "Evolving artificial neural networks." Proceedings of the IEEE **87**(9): 1423-1447.
- Xu, R. and D. Wunsch, II (2005). "Survey of clustering algorithms." Neural Networks, IEEE Transactions on **16**(3): 645-678.